Proving Correctness of Parallel Implementations of Transition System Specifications

Frank S. de Boer[1], Einar Broch Johnsen[2], Violet Ka I Pun[3], and S. Lizeth Tapia Tarifa[2]

[1] CWI, The Netherlands
F.S.de.Boer@cwi.nl
[2] Department of Informatics, University of Oslo, Norway
{einarj,sltarifa}@ifi.uio.no
[3] Department of Computing, Western Norway University of Applied Sciences, Norway
Violet.Ka.I.Pun@hvl.no

Abstract. The overall problem addressed in this paper is the long-standing problem of program correctness, and in particular programs that describe systems of parallel executing processes. We propose a new method for proving correctness of parallel implementations of high-level transition system specifications. The implementation language underlying the method is based on the model of active (or concurrent) objects. The method defines correctness in terms of a simulation relation between the transition system which specifies the program semantics and the transition system that is described by the correctness specification. The simulation relation itself abstracts from the fine-grained interleaving of parallel processes by exploiting a global confluence property of the particular model of active objects considered in this paper.

As a proof-of-concept we apply our method to the correctness of a parallel simulator of multicore memory systems.

1 Introduction

A long-standing challenge in Computer Science is the formal specification and verification of programs, notably that of parallel programs, e.g., multi-threaded Java programs which give rise to complex fine-grained interleaving of the parallel executing processes and interaction (locking) mechanisms.

Roughly, we can distinguish between logic and semantic based methods for establishing program correctness. Methods that are based on logic use assertions to express behavioral properties and generate proof conditions for their validation, which are usually discharged by interactive theorem proving. These methods are applicable to infinite-state systems and to actual programs used in practice (see for example [19] for the verification of a corrected version of the TimSort sorting program of the Java Collections Framework). One of the main complexities of the use of logics is due to the complexity of the specification of invariant properties and the interactive use of a theorem prover. On the other hand, model-checking is based on an automated state-space exploration of the program execution. This in general is restricted to finite-state systems and requires suitable abstraction techniques to master the state-space explosion problem.

The main contribution of this paper is a new method which supports the specification of an abstraction of the overall behavior of a parallel program in terms of a *transition system specification* (TSS, for short) [25,13,26]. Verifying that a parallel program satisfies such a correctness specification then involves establishing a *simulation relation* between the transition system describing the semantics of the parallel program and the system described by the TSS.

Our method supports a general approach to proving the correctness of parallel programs in two steps:

1. Verify global behavioral properties using a high-level formal model which abstracts from the complexity of the concurrency model of the target language to support inductive proofs of global properties.
2. Justify the correctness of the parallel implementation in the target language with respect to the high-level model in terms of a simulation relation.

Transition system specifications allow for the formal description of overall system behavior in a syntax-oriented, compositional way, using inference rules for local transitions and their composition. Process synchronization can be expressed abstractly using, e.g., conditions on system states and reachability conditions over transition relations as premises, and label synchronization for parallel transitions. This high level of abstraction greatly simplifies the verification of system properties. Whereas TSS is well-known as a formalism to define language semantics and reason about language meta-theory, it is also well-suited to describe specific systems in order to reason about, e.g., reachability or state invariance.

For the second step, we need an implementation language with a formal semantics (e.g., formalized by a TSS) which enables a simulation relation to be formally established. In this paper, we have opted for the *active object* language ABS [27,48] (ABS stands for *Abstract Behavioral Specification*). The semantics of the ABS language is formally defined by a TSS [29] and implemented by backends[4] in Erlang, Haskell, and Java, all of which support parallel execution. It has been developed and applied in the context of various EU projects, e.g., in the EU FP7 projects HATS[5] (*Highly Adaptable and Trustworthy Software using Formal Models*) and ENVISAGE[6] (*Engineering Virtualized Services*). In these projects, ABS has been extended and successfully applied to the formal modelling and analysis of software product families [17] and software services deployed on the Cloud [31]. The ABS tool suite [20,6,35,2,32,4,21,34] has been further applied to case studies, targeting cloud-based frameworks [53,40,30,39,3], railway operations [33] and computational biology.[7]

The parallel execution of active objects (see [18] for a survey of active object languages) is a direct consequence of decoupling method execution from method invocation by means of *asynchronous* method invocations. The ABS language further integrates a strict *encapsulation* of the local state of an active object with explicit language constructs for the *cooperative scheduling* of its method executions. Since the ABS language is tailored to the description of distributed systems, it abstracts from the order in which method invocations are generated.

In the definition of the simulation relation, cooperative scheduling allows the interleaving of methods in an active object to match the granularity of the transition rules of the corresponding TSS. Moreover, the parallel execution of active objects in ABS satisfies a *global confluence* property which allows to express *locally* the proof conditions of the simulation relation in a syntax-directed manner, abstracting from the fine-grained interleaving of the method executions.

As a proof-of-concept we introduce our method by its application to a parallel simulator of multicore memory systems. These memory systems generally use caches to avoid bottlenecks in data access from main memory, but caches introduce data duplication and require protocols to ensure coherence. Although data duplication is usually not visible to the programmers, the way a program interacts with these copies largely affects performance by moving data around to maintain coherence. To develop, test and optimize software for multicore architectures, we need correct, executable models of the underlying memory systems. A TSS of multicore memory systems with correctness proofs for cache coherency has been described in [9,10], together with a prototype implementation in the rewriting logic system Maude [15]. However, this fairly direct implementation of the TSS is not well suited to simulate large systems. Therefore we introduce in this paper a parallel implementation based on the active object model of the ABS language and apply our method to a proof of its correctness.

This paper extends [7] which describes a first version of the use case. The extension consists of a formalization of the novel idea of annotating ABS programs with TSS rules and the use of a global confluence property of the ABS semantics in the formal semantics (and verification) of these annotations. Because of the absence of this high-level specification of the simulation relation between the ABS program and the TSS, the ABS implementation in [7] has been developed largely independent of the TSS, which considerably complicated the correctness proof. In constrast, the application of our new methodology lead to a major refactoring of the ABS implementation described in [7], reflecting a correctness-by-design development methodology.

Plan of the paper In the following section we introduce the main concepts of the ABS language and in Section 3 the use of transition rules as annotations of ABS programs. In Section 4 we introduce the runtime syntax of the multicore TSS and in Section 5 we discuss its ABS implementation. Section 6 then introduces the correctness proof. Related work is discussed in Section 7 and in Section 8 some general conclusions are drawn and future work discussed.

2   ABS: Actors with Cooperative Concurrency

ABS is a modeling language for designing, verifying, and executing concurrent software [27,48]. The core language combines the syntax and object-oriented style of Java with the Actor model of concurrency [28], resulting in active objects which decouple communication and synchronization using asynchronous method calls and cooperative scheduling [18]. Asynchronous method calls generate processes (which execute the called methods) within the called (active) object and do not impose any synchronization between caller and callee. Instead, synchronization between different objects happens using (implicit) futures, with which the caller and callee may

---

Instruction	Meaning
new C	Creation of an instance of class C
switch (e){p_1=> $s_1 \cdots p_n$=> s_n}	Pattern matching
await b	Suspension on a Boolean condition
await e!m($e_1, \ldots, e_n$)	Suspension on termination of a asynchronous call
e!m($e_1, \ldots, e_n$)	Non-blocking asynchronous call
e.m($e_1, \ldots, e_n$)	Blocking synchronous call
m($e_1, \ldots, e_n$)	Inlined (recursive) self- call

Table 1. Basic ABS instructions used in this paper. Here, b is a Boolean expression, e and e_i denote expressions.

synchronise independently, at different times. Synchronization between different processes within an object is captured using cooperative scheduling. A process allows another process to be scheduled by means of explicit suspension points; rescheduling at the suspension point may depend on the resolution of a future or on a Boolean conditional. This mechanism allows the interleaving of different processes to be captured very precisely in ABS.

The imperative layer of synchronization and communication is complemented by a functional layer, used for computations over the internal data of objects. The functional layer combines parametric algebraic datatypes (ADTs) and a simple functional language with case distinction and pattern matching. ABS includes a library with predefined (Int, Bool, etc.) and parametric datatypes (lists, sets, maps, etc.) All other types and functions are user-defined.

In the following, the basic ABS instructions used in this paper (and shown in Figure 1) are explained in terms of some general synchronization patterns.

2.1 Synchronization Patterns

We discuss encodings in ABS of a basic locking mechanism, atomic operations, and a broadcast mechanism for global synchronization (using barriers).

Locks The basic mechanisms of asynchronous method calls and cooperative scheduling in ABS can be explained by the simple code example of a class Lock (Figure 1). It uses an **await** statement on a Boolean condition to model a binary semaphore, which enforces exclusive access to a common resource "lock", modeled as an instance of the class Lock (dynamically created by the execution of the expression **new** Lock). More specifically, execution of the take_lock method will be suspended by the **await** unlocked statement. This statement *releases the control*, allowing the scheduling of other (enabled) processes within the Lock object. When the local condition unlocked inside the Lock object has become true, the generated take_lock processes within the Lock object will compete for execution. The scheduled process then will terminate and return by setting unlocked to False.

```
class Lock {
  Bool unlocked = True;

  Unit take_lock{
    await unlocked;
    unlocked = False;
  }

  Unit release_lock{
    unlocked = True;
  }
}
```

Fig. 1. Lock implementation in ABS using await on Booleans.

In general, the *suspension points* defined by **await** statements define the granularity of interleaving of the processes of an object. The statement **await** lock!take_lock() will only suspend the process that issued the call (and release control in the caller object) until take_lock has returned. In contrast, a *synchronous* call lock.take_lock() in ABS will generate a process for the execution of the take_lock() method by the lock object and block (all the processes of) the caller object until the method returns.

Atomic operations The interleaving model of concurrency of ABS allows for a simple and high-level implementation of atomic operations. For example, Figure 2 shows a general ABS implementation of test-and-set instructions [5], where the concurrency model guarantees that the local */*test(input)*/* and */*set*/* instructions, assuming that they do not involve suspension points, are not interleaved and thus can be thought of as executed in a single atomic operation. In ABS test instructions can be implemented using the **switch**-instruction, which evaluates an expression that matches the resulting value

```
Bool TestandSet (/*input*/){
  Bool fail = False;
  switch /*test(input)*/ {
    True => /*set*/;
    False => fail = True; }
  return fail; }
```

Fig. 2. Test and set pattern in ABS.

against a pattern p in the different branches. This instruction has been mainly used to pattern match the ADTs used in the ABS program discussed in this paper. In the simplest case, this pattern can be replaced by an **if**-**else**-instruction. Instances of this atomic pattern can be observed in Figures 11 and 14, in the methods remove_inv and swap.

Broadcast synchronization Figure 3a shows how broadcast synchronization in a labelled TSS can be enforced simply by matching labels (an example is detailed in Section 4), thus abstracting from the implementation details of the implicit multi-party synchroniser. On the other hand, in programming langues like ABS the multi-party label synchronization needs to be programmed explicitly; Figure 3b illustrates the architecture of the ABS implementation in Figure 4.

a) Broadcast synchronization in a labelled TSS.

b) Broadcast synchronisation using an explicit synchroniser and barriers in the ABS model.
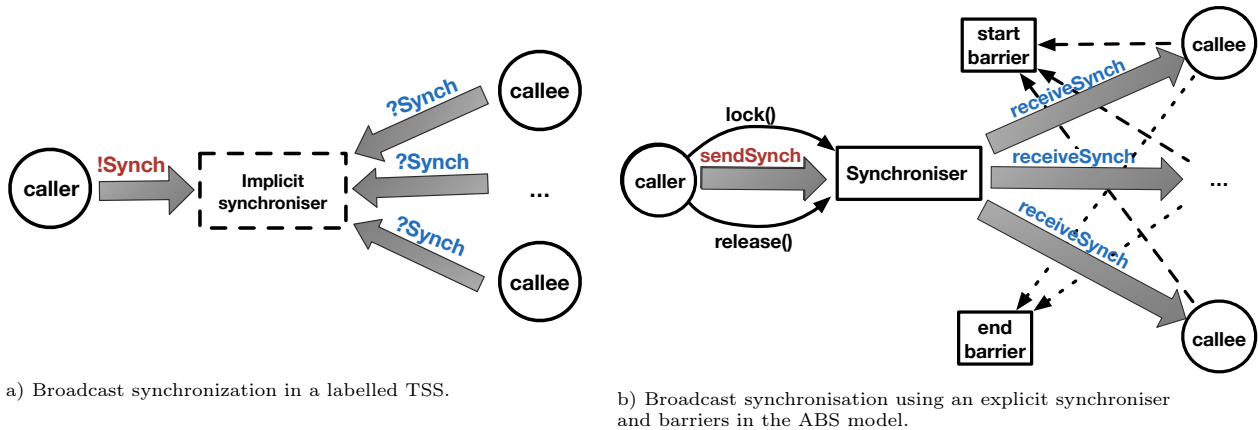
Fig. 3. Broadcast synchronisation patterns in TSS and ABS.

The class Broadcast serves as a template (or design pattern) for the implementation of a broadcast mechanism between its instances which is specified by the interface IBroadcast. The broadcastSync method encapsulates a synchronisation protocol between Broadcast instances which uses the additional classes Synchroniser and Barrier. This protocol consists of a synchronous call to the method sendSync of an instance of the class Synchronise (denoted by sync) which in turn asynchronously calls the method receiveSync of the objects stored in the set network of Broadcast instances, excluding the caller object executing the broadcastSync method. We abstract from whether the sync object is passed as parameter of the broadcastSync method or part of the local state of any Broadcast instance. The local computation specified by the receiveSync method by the objects in receivers is synchronized by calls of the method synchronise of the new instances start and end of class Barrier. That is, execution of this method by the start and end barriers synchronise the start and the termination of the execution of the method receiveSync by the objects in receivers and termination of the sendSync method itself. This is achieved by a "countdown" of the number of objects in receivers that have called the synchronise method plus one, in case of the end barrier. The synchronise method of the start barrier is called asynchronously (Line 15) and introduces a release point in order to avoid a deadlock that may arise when an object that has not yet called the synchronize method of the start barrier is blocked on a synchronous method call to an object that has already invoked (synchronously) the synchronize method of the start barrier. On the other hand, the corresponding call to the end barrier is synchronous to ensure that all the objects in receivers have completed their local computations. The additional synchronisation of the synchroniser object on the end barrier ensures that also the caller of the sendSync method is blocked until all the local computations specified by the receiveSync method have been completed.

Objects in ABS are input enabled, so it is always possible to call a method on an object. In our implementation, this scheme could give rise to inconsistent states if several objects start the protocol in parallel. To ensure exclusive access to the synchroniser at the start of the protocol, we add a lock to the synchroniser protocol, such that the caller must take the lock before calling sendSync and release the lock upon completion of the call. The resulting exclusive access to the synchroniser guarantees that its message pool contains at most one call to the method sendSync.

```
1   Interface IBroadcast {
2   Bool broadcastSync(...);
3   Unit receiveSync (IBarrier start, IBarrier end, ...)}
4
5   Class Broadcast implements IBroadcast, ...{
6   Bool broadcastSync(...){
7       Bool signal=False;
8       await sync!lock();
9       if /*test*/ { sync.sendSync(this,...); /*set*/; signal=True; }
10      sync.release();
11      return signal
12  }
13
14  Unit receiveSync(IBarrier start, IBarrier end, ...) {
15       await start!synchronise();
16      /*some local computation*/;
17      end.synchronise(); }
18  ...
19  }
20
21  Class Synchroniser (Set<IBroadcast> network) implements ISynchroniser {
22  Bool unlocked = True;
23  Unit lock(){ await unlocked; unlocked = False; }
24  Unit release(){unlocked = True; }
25  Unit sendSync(IBroadcast caller,...) {
26      Set<IBroadcast> receivers = remove(network,caller);
27      Int nrrecs= size(receivers);
28      IBarrier start = new Barrier(nrrecs);
29      IBarrier end = new Barrier(nrrecs+1);
30      foreach (receiver in receivers) { receiver!receiveSync(start,end,...); }
31      end.synchronise();}
32   }
33
34
35  class Barrier(Int participants) implements IBarrier {
36    Unit synchronise() { participants = participants − 1; await (participants == 0); }
37  }
```

Fig. 4. Global synchronisation pattern in ABS.

2.2 Semantics

ABS is a formally defined language [29]; in fact, its (operational) semantics is defined by a TSS which allows us to reason formally about the execution of ABS programs. The semantics of an ABS model can be described by a transition relation between global configurations. A global configuration is a (finite) set of object configurations. An object configuration is a tuple of the form $\langle oid, \sigma, p, Q \rangle$, where oid denotes the unique identity of the object, σ assigns values to the instance variables (fields) of the object, p denotes the currently executing process, and Q denotes a set of (suspended) processes (the object's "queue"). A process is a closure (τ, S) consisting of an assignment τ of values to the local variables of the statement S. We refer to [29] for the details of the TSS for deriving transitions $G \to G'$ between global configurations in ABS.

Although only one thread of control can execute in an active object at any time, cooperative scheduling allows different threads to interleave at explicitly declared points in the code, i.e., the **await** statements. When the currently executing process is suspended by an **await** statement, another (enabled) process is scheduled. Access to an object's fields is protected: any non-local (outside of the object) read or write to fields happens explicitly via method calls so as to mitigate race-conditions or the need for extensive use of explicit mutual exclusion mechanisms (locks).

Since active objects only interact via method calls and processes are scheduled non-deterministically, which provides an abstraction from the order in which the processes are generated by method calls, the ABS semantics

satisfies the following global confluence property (see also [7,54]) that allows to commute consecutive local computations steps of processes which belong to *different* objects.

Theorem 1 (Global confluence). *For any two transitions $G_1 \to G_2$ and $G_1 \to G_3$ that describe execution steps of processes of different objects, there exists a global configuration G_4 such that $G_2 \to G_4$ and $G_3 \to G_4$.*

An important consequence of the above global confluence property, which underlies the main results of this paper, is that we can restrict the global interleaving between processes by reordering the execution steps in an ABS computation. In particular, we can restrict the interleaving semantics of the ABS model taking into account general semantic properties of synchronous communication, and the implementation of locks and broadcast synchronization in ABS, as explained next.

Since a synchronous call of a method of *another* object in ABS, blocks all processes of the caller (object), the global confluence property allows further restricting the interleaving of the ABS processes so that the caller process is resumed *immediately* after the synchronous method invocation has terminated.

It is worthwhile to note that in general we can *not* assume that a method that is called synchronously in ABS is also scheduled *immediately for execution* because this would discard execution of other processes by the callee.

The global confluence property also allows abstracting from the internal computation steps of the above ABS implementation of the global (broadcast) synchronization pattern because it allows scheduling the processes generated by the *broadcast* method so that execution of this method is not interleaved with any other processes.

We can formalize the above in terms of the following notion of *stable* configurations.

Definition 1 (Stable configurations). *An object configuration is stable if the statement to be executed denotes the termination of an asynchronously called method (we assume a special runtime syntax which denotes such termination), or it starts with a synchronous call to another object or a* **await** *statement. A global ABS configuration is stable if all its object configurations are stable.*

Note that since synchronous self-calls are executed by inlining they do not represent an interleaving point.

In the sequel $G \Rightarrow G'$ denotes the transition relation which describes execution starting from a global stable configuration G to a next one G' (without passing intermediate global stable configurations). We distinguish the following three cases:

1. The transition $G \Rightarrow G'$ describes the *local* execution of a method by a single object.
2. The transition $G \Rightarrow G'$ describes the *rendez-vous* between the caller and callee of a synchronous method call in terms of the terminating execution of the called method, *followed* by the resumption of the suspended call.
3. The transition $G \Rightarrow G'$ describes the effect of executing the broadcast method, which thus describes the *global* synchronization of different objects.

This coarse-grained interleaving semantics of ABS forms the basis for the general methodology to prove correctness of ABS implementations of TSS specifications, described next.

3 The General Methodology

3.1 Annotating ABS with TSS Rules

For a general introduction of transition system specifications we refer to [25]. The general methodology for the development of ABS implementations of abstract TSSs is based on the coarse-grained interleaving described in Section 2 (denoted by the transition relation $\Rightarrow$): it allows focusing on the design of *local, sequential* code that implements the individual transition rules. This is reflected by the following use of transition rules as a *specification formalism* of ABS code. A *conditional transition rule* $b : R$ consists of a local Boolean condition b in ABS and a name R of a transition rule. We use sequences $b_1 : R_1; \ldots; b_n : R_n$ of conditional transition rules to annotate *stable points*. A stable point of a method definition denotes either its body or a substatement of its body that starts with an external synchronous call or an **await** statement. The idea is that each b_i is evaluated as a condition which identifies a *path* leading from the annotated stable point to a next one or to termination. The execution of this path should correspond to the application of the associated transition rule R_i. This correspondence involves a simulation relation, described below.

A sequence $b_1 : R_1; \ldots; b_n : R_n$ of conditional transition rules is evaluated from left to right, that is, the first transition rule from the left, the Boolean condition of which evaluates to true, is applicable. The case that all

Boolean conditions are false means that there does not exist a transition rule for *any* path from the annotated stable point to a next one or to termination (in the simulation relation all these paths would correspond to a "silent" transition). As a special case, we stipulate that for *any* path leading from a stable point *which has no associated annotation* to a next stable point (or to termination) there does *not* exist a corresponding transition rule. The use of annotations in the ABS code of the multicore memory system is shown in Section 5.2.

3.2   Correctness of the Implementation

The correctness of the ABS implementation with respect to a given TSS can be established by means of a simulation relation between the transition system describing the semantics of the ABS implementation and the transition system describing the TSS. The annotation of ABS code with (conditional) TSS rules provides a high-level description of the simulation relation, describing which rule(s) correspond with the execution of the ABS code from one stable point to a next one (or to termination). Underlying this high-level description, we define a simulation relation between ABS configurations and the runtime states of the TSS. This simulation relation is defined as an abstraction function α which maps every stable global ABS configuration G to a behavioral equivalent TSS configuration $\alpha(G)$ (see Section 6).

We restrict the simulation relation to *reachable* ABS configurations. A configuration G of the ABS program is reachable if $G_0 \Rightarrow^* G$, for some *initial* configuration G_0. In an initial configuration of the ABS multicore program all process queues are empty, and the only active processes are those about to execute the run methods of the cores. This restriction allows to use some general properties of the ABS semantics; e.g., upon return of a synchronous call, the local state of the calling object has not changed.

We can now express that a ABS program is a correct implementation of a TSS specification by proving that the following theorem holds, given an abstraction function α:

Definition 2 (Correctness). *Given an ABS program and a TSS, let α be an abstraction function from configurations of the ABS program to TSS configurations. The ABS program is a correct implementation of the TSS, if for any reachable configuration G and transition $G \Rightarrow G'$ of the ABS program we have that $\alpha(G) = \alpha(G')$ or $\alpha(G) \rightarrow \alpha(G')$.*

Because of the general confluence property of the ABS semantics to prove that α is a simulation relation, it suffices to verify the annotations of methods in terms of the abstraction function α. The general idea is that for each transition $G \Rightarrow G'$ which results from the execution from one stable point to a next one (or to termination), we have to show that $\alpha(G')$ results from $\alpha(G)$ by application of the enabled TSS rule associated with the initial stable point. In case no TSS rule is enabled, we have a "silent" step, that is, $\alpha(G) = \alpha(G')$.

4   A TSS for Multicore Memory Systems

Design decisions for programs running on top of a multicore memory system can be explored using simulators (e.g., [11,14,41,44]). Bijo et al. developed a TSS for multicore memory systems [9,10]. Taking this TSS as a starting point, we will study how a parallel simulator can be developed which implements the TSS and use this development to discuss the details of our methodology. We first introduce the main concepts of multicore memory systems and then look at their formalization in terms of a TSS.

A multicore memory system consists of cores that contain *tasks* to be executed, the *data layout* in main memory (indicating where data is allocated), and a system *architecture* consisting of cores with private multi-level caches and shared memory (see Figure 5). Such a system is parametric in the number of cores, the number and size of caches, and the associativity and replacement policy. Data is organized in blocks that move between the caches and the main memory.

Fig. 5. Abstract model of a multicore memory system.

For simplicity, we abstract from the data content of the memory blocks, assume that the size of cache lines and

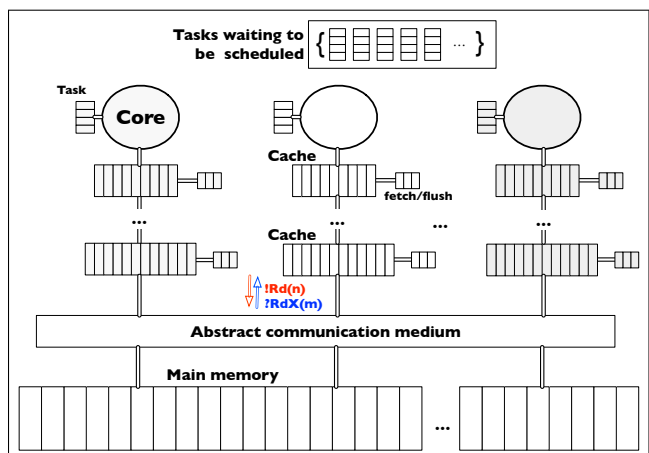| *Syntactic*		
categories.	*Definitions.*	
$cid \in CoreId$	$cf \in Config$	$::= \langle \overline{CR}, \overline{Ca}, M \rangle$
$caid \in CacheId$	$CR \in Core$	$::= cid \bullet rst$
$n \in Address$	$Ca \in Cache$	$::= caid \bullet M \bullet dst$
	$st \in Status$	$::= \{mo, sh, inv\}$
	$dap \in AccessPtns$	$::= \varepsilon \mid dap; dap \mid$ **read**$(n) \mid$ **write**(n)
	$rst \in RunLang$	$::= dap \mid rst; rst \mid$ **readBl**$(n) \mid$ **writeBl**(n)
	$dst \in DataLang$	$::= \varepsilon \mid dst + dst \mid$ **fetch**$(n) \mid$ **fetchBl**(n)
		$\mid$ **fetchW**$(n, n') \mid$ **flush**(n)

Fig. 6. Syntax of runtime configurations, where over-bar denotes sets (e.g., $\overline{CR}$).

memory blocks in main memory coincide and that a local reference to a memory block is represented directly by the corresponding memory address, and transfer memory blocks from the caches of one core to the caches of another core via the main memory. As a consequence, the tasks executed in the cores are represented as data access patterns, abstracting from their computational content.

Task execution on a core requires memory blocks to be transferred from the main memory to the closest cache. Each cache has a pool of instructions to move memory blocks among caches and between caches and main memory. Memory blocks may exist in multiple copies in the memory system. Consistency between different copies of a memory block is ensured using the standard cache coherence protocol MSI (e.g., [51]), with which a cache line can be either modified, shared or invalid. A *modified* cache line has the most recent value of the memory block, therefore all other copies are *invalid* (including the one in main memory). A *shared* cache line indicates that all copies of the block are consistent. The protocol's messages are broadcasted to the cores. The details of the broadcast (e.g., on a mesh or a ring) can be abstracted into an *abstract communication medium*. Following standard nomenclature, *Rd* messages request *read* access and *RdX* messages *read exclusive* access to a memory block. The latter invalidates other copies of the same block in other caches to provide write access.

We summarize the operational aspects of cache coherency with the MSI protocol. To access data from a memory block n, a core looks for n in its local caches. If n is not found in shared or modified state, a *read request* $!Rd(n)$ is broadcasted to the other cores and to main memory. The cache can *fetch* the block when it is available in main memory. Eviction is required if the cache is full, removing another memory block to free space. Writing to block n requires n to be in shared or modified state in the local cache; if it is in shared state, an *invalidation request* $!RdX(n)$ is broadcasted to obtain exclusive access. If a cache with block n in modified state receives a read request $?Rd(n)$, it *flushes* the block to main memory; if a cache with block n in shared state receives an invalidation request $?RdX(n)$, the cache line will be *invalidated*; the requests are discarded otherwise. Read and invalidation requests are broadcasted instantaneously in the abstract model, reflecting that signalling on the communication medium is orders of magnitude faster than moving data to or from main memory.

4.1   A TSS of Multicore Memory Systems

The multicore TSS describes the interactions between a core, caches, and the main memory. It further includes labeled transitions to model instantaneous broadcast. In general a model of the multicore TSS is a *transition system*. We refer to a model of the multicore TSS, which is parametric in the number of cores and caches, also as a Multicore Memory System (MMS, for short). The multicore TSS [9,10] is shown to guarantee correctness properties for data consistency and cache coherence (see, e.g., [16,52]), including the preservation of program order in each core, the absence of data races, and that stale data is never accessed.

We outline the main aspects of a simplified version of the multicore TSS which allows focusing on the main challenges of a correct distributed implementation. The runtime syntax is given in Figure 6. A configuration cf is a tuple consisting of a main memory M, cores $\overline{CR}$, caches $\overline{Ca}$ (we abstract from the tasks to be scheduled). A core ($cid \bullet rst$) with identifier cid executes *runtime statements rst*. A cache ($caid \bullet M \bullet dst$) with identifier $caid$ has a local cache memory M and data instructions dst. We assume that the cache identifier $caid$ encodes the cid of the core to which the cache belongs and its level in the cache hierarchy. We use $Status_\perp$ to denote the extension of the set $\{mo, sh, inv\}$ of status tags with the undefined value $\perp$. Thus, a memory $M : Address \rightarrow Status_\perp$ maps addresses n to either a status tag st or to $\perp$ if the memory block with address n is not found in M.

Data access patterns dap model tasks consisting of finite sequences of **read**(n) and **write**(n) operations to address n (that is, we abstract from control flow operations for sequential composition, non-deterministic choice, repetition, and task creation). The empty access pattern is denoted ε. Cores execute runtime statements rst,

which extend *dap* with **readBl**(n) and **writeBl**(n) to block execution while waiting for data. Caches execute *data instructions* from a multiset *dst* to fetch or flush a memory block with address n; here, **fetch**(n) fetches a memory block with address n, **fetchBl**(n) blocks execution while waiting for data, **fetchW**(n, n') waits for a memory block n' to be flushed before fetching n (this is needed when the cache is full), and **flush**(n) flushes a memory block.

The connection between the main memory and the caches of the different cores is modelled by an *abstract communication medium* which allows messages from one cache to be transmitted to the other caches and to main memory in a parallel instantaneous broadcast. Communication in the abstract communication medium is captured in the TSS by label matching on transitions. For any address n, an output of the form $!Rd(n)$ or $!RdX(n)$ is broadcasted and matched by its dual of the form $?Rd(n)$ or $?RdX(n)$. The syntax of the model is further detailed in [9,10]. For a complete overview of the transition rules we refer to A. In the next section, we will introduce these rules incrementally when discussing their ABS implementation.

The following auxiliary functions are used in the transition rules, given a cache identifier *caid*:

- $cid(caid)$ returns the identifier of the core to which the cache belongs;
- $lid(caid)$ gives the level at which the cache is located in the memory hierarchy;
- $first(caid)$ is true when $lid(caid) = 1$, otherwise false;
- $last(caid)$ is true when $lid(caid) = l$ where l is the number of caches in the hierarchy, otherwise false;
- $status(M, n)$ returns the status of block n in memory M or $\perp$ if the block is not found in M; and
- $select(M, n)$ determines the address where a block n should be placed in the cache, based on a cache associativity (e.g., random, set associativity or direct map) and a replacement policy (e.g., random or LRU).

5   The ABS Model of the Multicore Memory System

This section describes the translation of the multicore TSS into a model in ABS [8]. We explain the structural and behavioural correspondence between these two models.

5.1   The Structural Correspondence

The runtime syntax of the multicore TSS is represented in ABS by classes, user-defined datatypes and type synonyms, outlined in Figures 7–9. An ABS configuration consists of class instances to reflect the cores with their corresponding cache hierarchies and the main memory. Object identifiers guarantee unique names and object references are used to capture how cores and caches are related. These references are encoded in a one-to-one correspondence with the naming scheme of the multicore TSS.

A core $cid \bullet rst$ in the multicore TSS corresponds to an instance of the class Core in ABS, where a field currentTask of type RstList (as defined in Figure 9) represents the current list of runtime statements . Each instance of the class Core further holds a reference to the first level cache. An important design decision we made is to represent the runtime statements rst (of a core in the multicore TSS) as an ADT (see Figure 9). A core in ABS then drives the simulation by processing these runtime statements which in general requires information about the first-level cache. Alternatively, a core in ABS could delegate the processing of its runtime statements by calling corresponding methods of the first-level cache. However, this latter approach complicates the required callbacks.

A cache $caid \bullet M \bullet dst$ in the multicore TSS corresponds to an instance of class Cache with a class parameter nextLevel which holds a reference to the next level cache and a field cacheMemory which models the cache's memory M (of type MemMap, Figure 9). The multiset dst of a cache's data instructions (see Figure 6) is represented by corresponding *processes* in the message pool of the cache object in ABS. If the value of nextLevel is Nothing, then the object represents the last level cache (in the multicore TSS, a predicate *last* is used to identify the last level).

In addition, the ABS implementation of the global synchronisation with labels $!Rd(n)$ and $!RdX(n)$ used in the multicore TSS is based on the global synchronisation pattern as described in Figure 4. However, instead of distinguishing between these two labels by means of an additional parameter, we introduce two corresponding broadcast interfaces: ¨
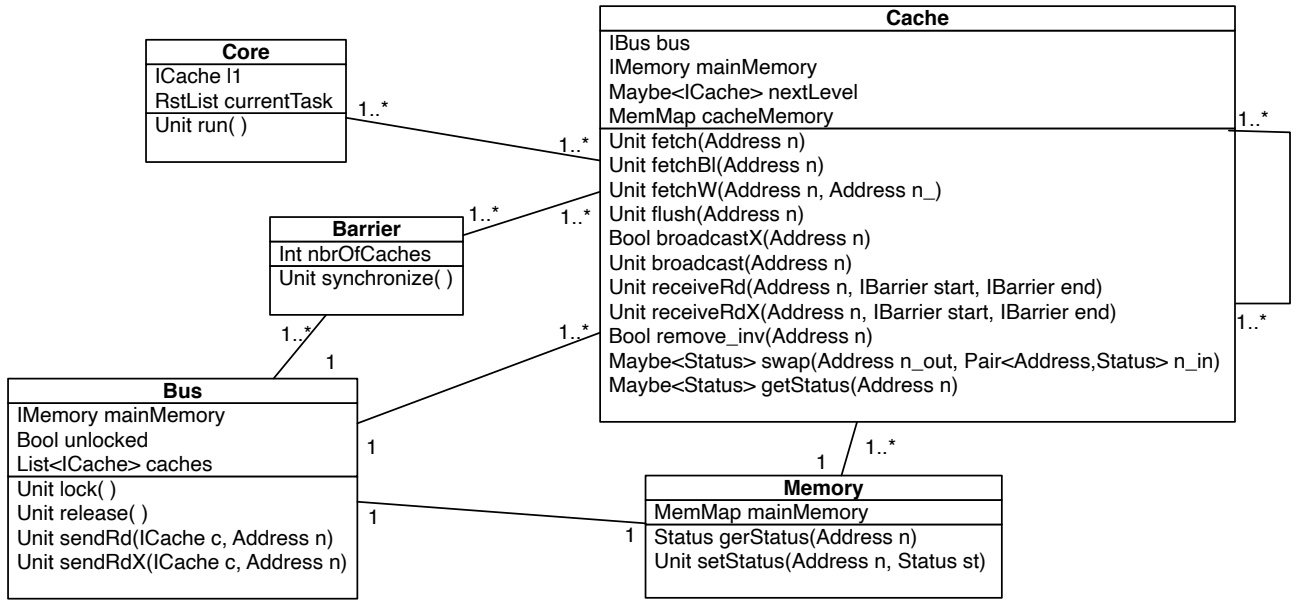

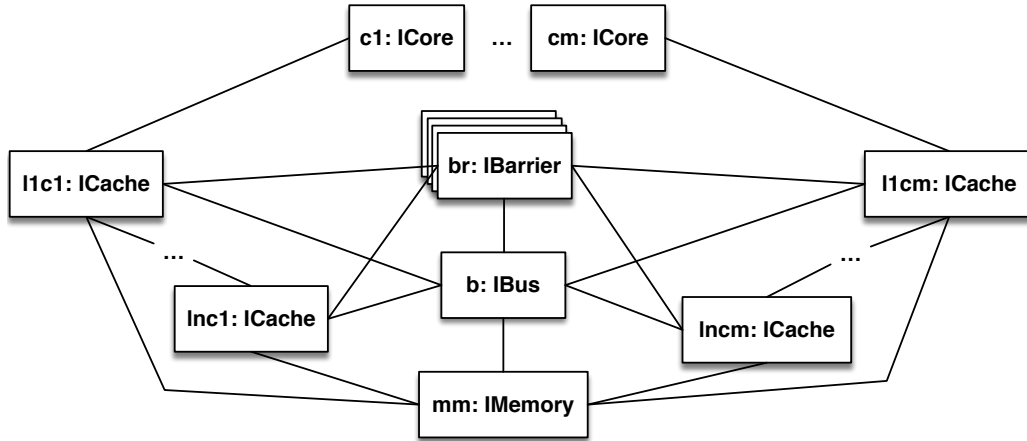**Fig. 7.** Class diagram of the ABS model.

Fig. 8. Object diagram of an initial configuration.

```
1   data Rst = Read(Address) | ReadBl(Address) | Write(Address) | WriteBl(Address);
2   data Status = Sh | Mo | In;
3   type RstList = List<Rst>;
4   type Address = Int;
5   type MemMap = Map<Address,Status>;
```

Fig. 9. Abstract data types of the model of the multicore memory system.

```
1   Interface IBroadcast {
2   Bool broadcast(...);
3   Unit receiveRd (IBarrier start, IBarrier end, ...)}
4
5   Interface IBroadcastX {
6   Bool broadcastX(...);
7   Unit receiveRdX (IBarrier start, IBarrier end, ...)}
```

```
1   Unit run() {
2     if currentTask!=Nil {
3     switch (currentTask) {
4        Cons(rst, rest) =>
5        switch (rst) {
6          Read(n) => {
7             removed = l1.remove_inv(n); // removed==True: PRRD₂; removed==False: PRRD₁
8             if (removed){
9                l1!fetch(n);
10               currentTask = Cons(ReadBl(n),rest); }
11             else {currentTask = rest; } }
12         ReadBl(n) => {
13            status = l1.getStatus(n); // status!=Nothing: PRRD₃
14          if (status != Nothing) currentTask = Cons(Read(n),rest); }
15        Write(n) => {
16          status = l1.getStatus(n); // status==Just(Mo): PRWR₁
17          switch (status) {
18            Just(Mo) => {currentTask = rest; }
19            Just(Sh) => {
20              Bool res = l1.broadcastX(n); // res==True: PRWR₂/SYNCHX
21              if (res) {currentTask = rest; } }
22            _ => { Bool removed = l1.remove_inv(n); // removed==True: PRWR₃
23                 if (removed){l1!fetch(n); currentTask = Cons(WriteBl(n),rest); } } } }
24        WriteBl(n) => {
25          Maybe<Status> status = l1.getStatus(n); // status!=Nothing: PRWR₄
26          if (status != Nothing)
27             currentTask = Cons(Write(n),rest); }
28       }
29     this ! run(); }
30  }
```

Fig. 10. Annotated run method.

The class Cache then provides an implementation of both interfaces following the template of the class Broadcast in Figure 4. The ABS class Bus, on the other hand, follows the template of the Synchroniser class with the two versions sendRd and sendRdX of the method sendSync.

The object diagram in Figure 8 shows an initial configuration corresponding to the one depicted in Figure 5.

5.2   The Behavioural Correspondence

We next discuss the ABS implementation of the transition rules of the multicore TSS, and the ABS synchronization patterns described in Section 2. We observe that the combination of *asynchronous method calls* and *cooperative scheduling* in ABS is crucial because of the *interleaving* inherent in the multicore TSS, which requires that objects are able to process other requests while executing a method in a controlled way; e.g., caches need to flush memory blocks while waiting for a fetch to succeed.

The Annotated ABS Multicore Implementation The classes Core and Cache pose the main implementation challenges. Here we explain the implementation of the run method (Figure 10) of the class Core (which is its only method) informally, in terms of its annotations (as introduced in Section 3.1). In Section 6 we introduce a formal semantics of these annotations as a high-level description of a simulation relation, and prove the correctness of the class Cache.

The run method may generate synchronous calls to the auxiliary methods in Figure 11. The method remove_inv instantiates the test-and-set pattern of Figure 2. The method broadcastX is an instance of the global synchronization pattern described in Section 2, Figure 4. The method sendRdX of the global synchroniser bus asynchronously calls the method receiveRdX, see Figure 12, of all caches (except for the calling cache), using the barrier synchronization described in Section 2.

Since the stable point at the beginning of the run method has no associated annotation, by definition (see Section 3.1), for *any* path from the beginning to a next stable point (or to termination) there does *not*

```
1   Maybe<Status> getStatus(Address n) { return lookup(cacheMemory,n); }

1   Bool remove_inv(Address n){
2     Bool answer = False;
3     switch (lookup(cacheMemory,n)) {
4       Nothing => { answer = True; }
5       Just(In) =>{ cacheMemory = removeKey(cacheMemory,n); answer = True; }
6       _ => skip; }
7     return answer; }

1   Bool broadcastX(Address n) {
2     Bool res = False;
3     await bus!lock(); //(lookup(cacheMemory,n) ==Just(Sh)): PRWR₂/SYNCH
4     if (lookup(cacheMemory,n) ==Just(Sh)){
5       bus.sendRdX(this, n);
6       cacheMemory = put(cacheMemory,n,Mo);
7       res = True; }
8       bus.release();
9     return res; }
```

Fig. 11. Methods getStatus, remove_inv, and broadcastX of class Cache.

```
1   Unit receiveRdX(Address n,IBarrier start,IBarrier end) {
2     // lookup(cacheMemory,n))==Just(Sh): INVALIDATE-ONE-LINE;
3     // lookup(cacheMemory,n))!=Just(Sh): IGNORE-INVALIDATE-ONE-LINE
4     await start!synchronize();
5     switch (lookup(cacheMemory,n)) {
6       Just(Sh) => {cacheMemory = put(cacheMemory,n,In); }
7       _ => skip;
8     }
9     end.synchronize(); }
```

Fig. 12. Annotated receiveRdX method.

correspond a transition rule (of the multicore TSS). For example, there is no transition rule corresponding to the case that the run method terminates when curentTask==Nil (note that because of the structural correspondence also the corresponding core has no runtime statements *rst* to execute). Similarly, there are no transition rules corresponding to the execution of the code from the beginning of the method to the synchronous calls to the auxiliary methods remove_inv (Line 7) and getStatus (Lines 13, 16, 25) of the first level cache which, besides the pattern matching, only consists of the call itself.

The condition of the annotation removed==True : PRRD_2 (Line 7) associated with the synchronous call to the remove_inv method describes the path which leads from its execution and return via the **then**-branch of the subsequent **if**-statement to the termination of the run method (after it has called itself again asynchronously). According to the annotation, the execution of this path corresponds to the PRRD_2 transition rule:

$$(\text{PRRD}_2)$$
$$\frac{first(caid) = true \quad cid(caid) = c \quad status(M,n) \in \{inv, \bot\}}{(c \bullet \text{ read}(n); rst \ ), \ (caid \bullet \ M \bullet dst \ ) \rightarrow}$$
$$(c \bullet \text{ readBl}(n); rst \ ), \ (caid \bullet \ M[n \mapsto \bot] \bullet dst + \text{fetch}(n) \ )$$

This rule handles the case when a core intends to read a memory block with address n, which is not found in the first level cache. The core will then be blocked while waiting for the memory block to be fetched either from the lower level caches or main memory. Note that the condition as returned by the remove_inv method signals that the status of the address of the first level cache is undefined or invalid.

On the other hand, the condition removed==False describes the path which leads from its execution and return via the **else**-branch (Line 11), which also leads to the termination of this invocation of the run method.

According to the annotation, the execution of this path corresponds to the $\textsc{PrRd}_1$ transition rule:

$$(\textsc{PrRd}_1)$$

$$\frac{first(caid) = true \quad cid(caid) = c \quad status(M, n) \in \{sh, mo\}}{(c \bullet \textbf{read}(n); rst\,),\ (caid \bullet M \bullet dst) \rightarrow (c \bullet rst\,),\ (caid \bullet M \bullet dst)}$$

This rule covers the case when the memory block to be read by a core is found in its first level cache. . Note that the condition as returned by the remove_inv method implies that the status of the address of the first level cache is either shared or modified.

Next we consider the annotation status != Nothing : $\textsc{PrRd}_3$ of the synchronous call to the getStatus method (Line 13). Its condition describes the execution path which leads from the execution and return of the called getStatus method to termination of the run method via the **then**-branch of the subsequent **if**-statement (Line 14). According to the annotation, the execution of this path corresponds to the $\textsc{PrRd}_3$ transition rule:

$$(\textsc{PrRd}_3)$$

$$\frac{first(caid) = true \quad cid(caid) = c \quad n \in dom(M)}{(c \bullet \textbf{readBl}(n); rst\,),\ (caid \bullet M \bullet dst) \rightarrow (c \bullet \textbf{read}(n); rst\,),\ (caid \bullet M \bullet dst)}$$

This rule unblocks the core from waiting when n (i.e., the block to be read) is found in the first level cache. On the other hand, there does not exist a transition rule which corresponds to the execution path described by the condition status==Nothing. This path leads from the execution of the called getStatus method directly to the termination of the run method without an update of the (local) state, e.g., currentTask is not updated. In other words, the evaluation of the readBl(n) instruction in ABS involves *busy waiting* until the status returned by the first level cache is defined. Alternatively, this could be implemented by calling synchronously a method of the first level cache which simply executes the statement **await** lookup(cacheMemory,n)!=Nothing.

The annotation of the synchronous call to the method getStatus (Line 25) involves the rule

$$(\textsc{PrWr}_4)$$

$$\frac{first(caid) = true \quad cid(caid) = c \quad n \in dom(M)}{(c \bullet \textbf{writeBl}(n); rst\,),\ (caid \bullet M \bullet dst) \rightarrow (c \bullet \textbf{write}(n); rst\,),\ (caid \bullet M \bullet dst)}$$

This annotation is explained in a similar manner as the annotation of the synchronous call to the getStatus method on Line 13. This rule unblocks the core from waiting when n (i.e., the block to be written) is found in the first level cache.

We consider next the annotation status==Just(Mo) : $\textsc{PrWr}_1$ of the synchronous call to the method getStatus (Line 16) . Its condition describes the execution path which leads from the execution of the called getStatus method and subsequent execution of the **switch** statement to termination of the run method. According to the annotation, the execution of this path corresponds to the $\textsc{PrWr}_1$ transition rule:

$$(\textsc{PrWr}_1)$$

$$\frac{first(caid) = true \quad cid(caid) = c \quad status(M, n) = mo}{(c \bullet \textbf{write}(n); rst\,),\ (caid \bullet M \bullet dst) \rightarrow (c \bullet rst\,),\ (caid \bullet M \bullet dst)}$$

This rule allows a core to write to memory block n if the block is found in a modified state in the first level cache. On the other hand, in case the condition does not hold, according to the annotation no transition rules correspond to the execution paths which lead from the execution of the called getStatus method to the next stable points, i.e., the synchronous calls to the methods broadcastX and remove_inv (Lines 20 and 22, respectively).

The condition of the annotation res==true : $\textsc{PrWr}_2/\textsc{SynchX}$ of the synchronous call to the broadcastX method (Line 20) of the first level cache describes the path which leads from the execution of the broadcastX method, followed by the execution of the subsequent if-statement to termination of the run method (after an update of currentTask and calling the run method again asynchronously). According to the annotation this path corresponds to the global synchronization rule

$$(\textsc{SynchX})$$

$$\frac{CR \notin \overline{CR_1} \quad CR, \overline{Ca} \xrightarrow{!RdX(n)} CR', \overline{Ca}'}{\langle \overline{CR_1} \cup \{CR\},\ \overline{Ca},\ M \rangle \rightarrow \langle \overline{CR_1} \cup \{CR'\},\ \overline{Ca}',\ M[n \mapsto inv] \rangle}$$

where the second premise is generated by successive applications of the rule

$$(\text{Synch-DistX})$$

$$\frac{Ca_1 \notin \overline{Ca} \quad CR, \overline{Ca} \xrightarrow{!RdX(n)} CR', \overline{Ca'} \quad Ca_1 \xrightarrow{?RdX(n)} Ca_2}{CR, \overline{Ca} \cup \{Ca_1\} \xrightarrow{!RdX(n)} CR', \overline{Ca'} \cup \{Ca_2\}}$$

This latter rule itself is triggered by the following rules

$$(\text{PrWr}_2)$$

$$\frac{first(caid) = true \quad cid(caid) = c \quad status(M, n) = sh}{(c \bullet \textbf{write}(n); rst\ ), (caid \bullet\ M\ \bullet dst) \xrightarrow{!RdX(n)} (c \bullet\ rst\ ), (caid \bullet\ M[n \mapsto mo]\ \bullet dst)}$$

$$(\text{Invalidate-One-Line})$$

$$\frac{status(M, n) = sh}{caid \bullet\ M\ \bullet dst \xrightarrow{?RdX(n)} caid \bullet\ M[n \mapsto inv]\ \bullet dst}$$

and

$$(\text{Ignore-Invalidate-One-Line})$$

$$\frac{status(M, n) \in \{inv, \bot\}}{caid \bullet M \bullet dst \xrightarrow{?RdX(n)} caid \bullet M \bullet dst}$$

These rules together capture the broadcast mechanism for invalidation in the multicore memory system. Rule PrWr_2 corresponds to the case where a core writes to a memory block n that is marked as shared in its first level cache, which requires broadcasting an invalidation message, $!RdX(n)$, to all the other caches. This is achieved by triggering the global synchronization rules SynchX and Synch-DistX. While the former identifies the core CR that broadcasts the invalidation message, the latter recursively propagates the message, $?RdX(n)$, to the other caches. Depending on the local status of memory block n in the recipient cache, the recipient cache will either invalidate the local copy of the block ($\text{Invalidate-One-Line}$), or ignore the message ($\text{Ignore-Invalidate-One-Line}$).

To explain this application of the SynchX rule, we have a closer look at the definition of the broadcastX method. Its body involves an instance of the global synchronization pattern (Figure 4). As discussed in Section 2, because of the global confluence property, we may assume that its execution is atomic, i.e., not interleaved with any process that it has not generated. The synchronous call to the sendRdX method of the bus generates asynchronous calls to the receiveRdX method (Figure 12) of all caches except the one that initiated the global bus synchronization. Following the general global synchronization pattern (Figure 4), these method calls are synchronized by a start and an end barrier. The two conditions of the annotation at the beginning of the receiveRdX method describe the two possible execution paths and their corresponding transition rules $\text{Invalidate-One-Line}$ and $\text{Ignore-Invalidate-One-Line}$.

In case the condition res==true does not hold, according to the annotation, no transition rule corresponds to the execution of the broadcastX method. In this case the bus synchronization, as invoked by the broadcastX method (Figure 11), failed because the status of the address of the first level cache is not shared anymore (as required by the PrWr_2 rule). Consequently, the processing of the write(n) instruction itself fails and it will be processed again by the asynchronous self call to the run method.

We conclude the informal explanation of the annotated run method with the annotation removed==True : PrWr_3 of the synchronous call to the method remove_inv (Line 22). Its condition describes the path that corresponds to the transition rule:

$$(\text{PrWr}_3)$$

$$\frac{first(caid) = true \quad cid(caid) = c \quad status(M, n) \in \{inv, \bot\}}{\begin{array}{c}(c \bullet \textbf{write}(n); rst\ ), (caid \bullet\ M \bullet dst\ ) \to \\ (c \bullet \textbf{writeBl}(n); rst\ ), (caid \bullet\ M[n \mapsto \bot]\ \bullet dst + \textbf{fetch}(n)\ )\end{array}}$$

This rule handles the case when a core tries to write to a memory block with address n, which is either invalid or not found in the first level cache. The core will then be blocked while the memory block is fetched from

the lower level cache or from the main memory. On the other hand, according to the annotation, no transition rule corresponds to the execution path that is described by the negation of the condition. Note that this covers the case when the status returned by getStatus (Line 16) has changed; i.e., the status of the memory block is no longer undefined or invalid. As above, the run method terminates without having successfully processed the write(n) task, which will be evaluated again by the next asynchronous invocation of the run method.

In the next section we show how to formally validate the annotations in terms of a simulation relation.

6   The Simulation Relation

We now establish the correctness of the ABS implementation of the multicore memory system with respect to the multicore TSS specification. First we observe that the structural correspondence described in Section 5 only relates the class diagram of the ABS program (Figure 7) and the syntax of the runtime configurations (Figure 6) of the multicore TSS. To relate *behavioral* information, we define the abstraction function α which maps every stable global ABS configuration G to a structurally equivalent configuration $\alpha(G)$ which additionally provides a *one-to-one mapping* between the *observable* processes of the instances of the ABS class Cache and the *dst* instructions of the corresponding TSS cache representation Ca, such that the actual address of the associated *dst* instruction equals the value of the corresponding formal parameter of the ABS process. The observable ABS processes are those that stem from an asynchronous call of a method that corresponds with a *dst* instruction (like fetch, etc.).

We have the following main theorem stating that the ABS multicore program is a correct implementation of the multicore TSS as an instance of Theorem 2 (recall that $\Rightarrow$ denotes the transition relation between stable ABS configurations):

Theorem 2. *Let G be a reachable stable global configuration of the ABS multicore model. If $G \Rightarrow G'$ then $\alpha(G) = \alpha(G')$ or $\alpha(G) \to \alpha(G')$.*

Proof of Theorem 2. Because of the general confluence property of the ABS semantics, it suffices to verify the annotations of the run method and the methods of the Cache class that correspond to the *dst* instructions in terms of the simulation relation α. Here we detail the analysis of the Cache class.

We first verify the annotations of the fetch method (Figure 13), identifying stable points by their line numbers. The fetch method involves synchronous calls to the auxiliary methods broadcast (Figure 13) and swap (Figure 14). The method broadcast describes an instance of the global synchronization pattern (Figure 4). The method sendRd of the bus asynchronously calls the method receiveRd, see Figure 15, of all caches (except for the calling cache), using the barrier synchronization (again, see Figure 4, Section 2). The swap method is an instance of the test-and-set pattern, shown in Figure 2.

Let $G \Rightarrow G'$ describe the execution of an invocation of the fetch method from one stable point to a next one (or to termination), and let $caid \bullet M \bullet dst + \mathbf{fetch}(n)$ be the cache in $\alpha(G)$ that corresponds to the cache in G executing the fetch method, where n denotes the value of the formal parameter n of the executing invocation of the fetch method. Further, let $caid' \bullet M' \bullet dst'$ be the cache in $\alpha(G)$ that corresponds to the next level cache in $\alpha(G)$, if defined. If such a next level cache exists, we have that $lid(caid') = lid(caid) + 1$ and $cid(caid) = cid(caid')$. By $\overline{M}$ we denote in the following the *main* memory in $\alpha(G)$. We have the following case analysis of $G \Rightarrow G'$ that describe the execution of an invocation of the fetch method from one stable point to a next one (or to termination).

Lines 1 $\Rightarrow$ 5. In this case, $G \Rightarrow G'$ involves the execution of the fetch method (by a cache object) starting from the beginning of the method (Line 1) to the synchronous call of the method remove_inv (Line 5) of the next level cache (note that thus nextCache != Nothing holds). Since this invalidates the path condition of the annotation associated with the beginning of the fetch method, by definition (see Section 3.1), there is *no* transition rule (of the multicore TSS) which corresponds to $G \Rightarrow G'$. This execution only adds this call to the queue of the next level cache and α abstracts from invocations of the method remove_inv, so it follows that $\alpha(G') = \alpha(G)$.

Lines 5 $\Rightarrow$ 8. In this case, $G \Rightarrow G'$ consists of the path which leads from the execution of the called remove_inv method via the **then**-branch of the subsequent **if**-statement to the termination of the fetch method (Line 8). According to the annotation removed==true : LC-MISS, this execution path should correspond to the

```
1   Unit fetch(Address n){
2   // nextLevel==Nothing: LLC-MISS/SYNCH
3     switch (nextLevel) {
4       Just(nextCache) => {
5         Bool removed = nextCache.remove_inv(n); // removed==true: LC-MISS
6         if (removed){
7           nextCache!fetch(n);
8           this!fetchBl(n); }
9         else { Pair<Address,Status> selected = select(cacheMemory, n);
10          Maybe<Status> s = nextCache.swap(n,selected);
11          // s!=Nothing\& fst(selected)==n: LC-HIT₂;
12          // s!=Nothing\& fst(selected)!=n: LC-HIT₁
13          if (s != Nothing){ if (fst(selected)!=n){
14            cacheMemory = removeKey(cacheMemory,fst(selected)); }
15            cacheMemory = put(cacheMemory, n,fromJust(s)); }
16          else this!fetch(n); } }
17        _ => { this.broadcast(n);
18              this!fetchBl(n); } } }
19
20  Unit broadcast(Address n){
21    await bus!lock();
22    bus.sendRd(this, n);
23    bus.release(); }
```

Fig. 13. The annotated fetch method.

following application of the LC-MISS rule:

$$(\text{LC-MISS})$$
$$\frac{lid(caid') = lid(caid) + 1 \quad cid(caid) = cid(caid') \quad status(M',n) \in \{inv, \bot\}}{(caid \bullet M \bullet \; dst + \mathbf{fetch}(n) \;), \; (caid' \bullet \; M' \bullet dst' \;) \rightarrow}$$
$$(caid \bullet M \bullet \; dst + \mathbf{fetchBl}(n) \;), \; (caid' \bullet \; M'[n \mapsto \bot] \bullet dst' + \mathbf{fetch}(n) \;)$$

This rule handles the situation where a cache is trying to fetch a memory block n from its next level cache, but the block is either invalidated or does not exist in the cache. The fetch-method in this cache will then be suspended; fetch is propagated to the next level cache and the memory block n will be removed from the next level cache. Since in this case the method remove_inv has returned the Boolean value "True" we can infer statically from its code (Figure 11, Section 5) that the initial status of the given address in the next level cache is Nothing or Just(In). Thus, the conditions for this application of the LC-MISS rule are enabled in $\alpha(G)$. Moreover, we can statically infer in this case that the execution of the fetch method is simulated by the updates $dst + \mathbf{fetchBl}(n)$ and $dst' + \mathbf{fetch}(n)$, and the remove_inv method is simulated by the update $M'[n \mapsto \bot]$. We conclude that $\alpha(G) \rightarrow \alpha(G')$ by this application of the LC-MISS rule.

Lines 5 ⇒ 10. In this case, $G \Rightarrow G'$ consists of the path which leads from the execution of the remove_inv method via the **else**-branch of the subsequent **if**-statement to the (synchronous) call of the swap method (Line 10). Since in this case the method remove_inv has returned the Boolean value "False", the configuration G' results from G by initializing the local variable selected and queuing the call of the swap method. Abstracting from the definition of the ABS select function, which picks a cache line for eviction to give space to a newly fetched memory block, we simply assume that the first element of the pair denoted by the ABS expression select(cacheMemory,n) equals the address $select(M,n)$ and its second element equals the status of this address.[9] By definition of the simulation relation which abstracts from local variables and invocations of auxiliary methods like the swap method, it follows that $\alpha(G') = \alpha(G)$.

Lines 10 ⇒ 15. In this case, $G \Rightarrow G'$ describes one of the two paths which lead from the execution of the swap method via the **then**-branch of the subsequent **if**-statement (so s != Nothing) to the termination of the fetch method (Line 15). We infer statically from s != Nothing that $M'(n) = s'$, where $s' = sh \vee s' = mo$. Further,

---

[9] In the actual ABS implementation, an extra parameter is used to capture the maximum size of the cache to check if there is free space to fetch a memory block n from its next level.

```
1   Maybe<Status> swap(Address n_out, Pair<Address,Status> n_in) {
2     Maybe<Status> tmp = Nothing;
3     switch (lookup(cacheMemory,n_out)) {
4       Nothing => skip;
5       Just(In) => skip;
6       _ => {
7         tmp = lookup(cacheMemory,n_out);
8         cacheMemory = removeKey(cacheMemory,n_out);
9         if (fst(n_in)!=n_out) {
10          cacheMemory = put(cacheMemory, fst(n_in), snd(m_in)); } } }
11    return tmp; }
```

Fig. 14. The swap method.

```
1   Unit receiveRd(Address n,IBarrier start,IBarrier end) {
2     // lookup(cacheMemory,n))==Just(Mo): Flush-One-Line;
3     // lookup(cacheMemory,n))!=Just(Mo): Ignore-Flush-One-Line
4     await start!synchronize();
5     switch (lookup(cacheMemory,n)) {
6       Just(Mo) => this!flush(n);
7       _ => skip;
8     }
9     end!synchronize(); }
```

Fig. 15. The annotated receiveRd method.

because G is reachable we infer from the semantics of synchronous calls that selected==select(cacheMemory,n) holds in G for the given cache object executing the fetch method. Finally, we observe that fst(selected) == n holds at Line 13 of the fetch method if and only if fst(n_in) == n_out holds at Line 9 of the swap method. There are two cases, depending on the value of fst(selected).

First, let fst(selected) != n, that is, $select(M, n) \neq n$. It follows that the enabling conditions of the LC-Hit$_1$ rule

$$(\text{LC-Hit}_1)$$
$$lid(caid') = lid(caid) + 1 \quad cid(caid) = cid(caid')$$
$$select(M, n) = m \quad n \neq m \quad M(m) = s \quad M'(n) = s' \quad s' = sh \vee s' = mo$$

$$(caid \bullet M \bullet dst + \textbf{fetch}(n)), \; (caid' \bullet M' \bullet dst') \rightarrow$$
$$(caid \bullet M[m \mapsto \bot, n \mapsto s'] \bullet dst), \; (caid' \bullet M'[n \mapsto \bot, m \mapsto s] \bullet dst')$$

are satisfied in $\alpha(G)$. Rule LC-Hit$_1$ addresses the case where a cache finds n, i.e., the block to be fetched, in its next level with status shared or modified, and there is no free space in the cache to place n. In order to fetch n, the rule selects a memory block m in the current cache to be swapped with n in the next level cache. According to the annotation s != Nothing & fst(selected) != n : LC-Hit$_1$, this rule should correspond to the path identified by its condition. Since fst(n_in) != n_out, the execution of the swap method by the next level cache in ABS is simulated by the update $M'[n \mapsto \bot, m \mapsto s]$ and, on the other hand, the execution of the two assignments (Lines 14 and 15 of the fetch method) is simulated by the update $M[m \mapsto \bot, n \mapsto s']$.

Next, let fst(selected)==n, that is, $select(M, n) == n$. It follows that the enabling conditions of the LC-Hit$_2$ rule

$$(\text{LC-Hit}_2)$$
$$lid(caid') = lid(caid) + 1 \quad cid(caid) = cid(caid')$$
$$select(M, n) = n \quad M'(n) = s' \quad s' = sh \vee s' = mo$$

$$(caid \bullet M \bullet dst + \textbf{fetch}(n)), \; (caid' \bullet M' \bullet dst') \rightarrow$$
$$(caid_i \bullet M[n \mapsto s'] \bullet dst), \; (caid' \bullet M'[n \mapsto \bot] \bullet dst')$$

are satisfied. This rule addresses the case where a cache finds n, i.e., the memory block to be fetched, in its next level cache with status shared or modified, and there is free space in the current cache to place n. According

to the annotation s != Nothing & fst(selected) == n : LC-Hit$_2$, this rule should correspond to the the path identified by its condition. Since fst(n_in) == n_out, in this case the execution of the swap method by the next level cache in ABS is simulated by the $M'[n \mapsto \bot]$ update and, on the other hand, the execution of the assignment (Line 15 of the fetch method) is simulated by the update $M[n \mapsto s']$.

Lines 10 ⇒ 16. In this case, $G \Rightarrow G'$ describes the execution path which leads from the return of the swap method via the **else**-branch of the subsequent **if**-statement (so s == Nothing) to the termination of the fetch method. We infer statically from the code of the swap method and the condition s == Nothing that the status of the address denoted by the formal parameter n of the fetch method of the next level cache is undefined or invalid. According to the annotation of the swap method, no rule of the multicore TSS is applicable. Since this invocation of the fetch method terminates after an asynchronous self-call transmitting the same address, we have that $\alpha(G) = \alpha(G')$.

Lines 1 ⇒ 18 In this case, $G \Rightarrow G'$ involves the execution of the fetch method starting from the beginning of the method which leads to termination after the execution of the broadcast method and the asynchronous self-call to the fetchBl method. The method broadcast implements an instance of the global synchronization pattern (Figure 4). The synchronous call to the sendRd method of the bus generates asynchronous calls to the receiveRd method (Figure 12) of all caches except the one that initiated the global bus synchronization. According to the annotation LLC-Miss/Synch, this execution path corresponds to the global synchronization rule

$$
\text{(Synch)}
$$
$$
\frac{\overline{Ca} \xrightarrow{\ !Rd(n)\ } \overline{Ca'}}{\langle \overline{CR},\ \overline{Ca},\ M \rangle \to \langle \overline{CR},\ \overline{Ca'},\ M \rangle}
$$

where the premise is generated by successive applications of the rule

$$
\text{(Synch-Dist)}
$$
$$
\frac{Ca_1 \notin \overline{Ca} \quad \overline{Ca} \xrightarrow{\ !Rd(n)\ } \overline{Ca'} \quad Ca_1 \xrightarrow{\ ?Rd(n)\ } Ca'_2}{\overline{Ca} \cup \{Ca_1\} \xrightarrow{\ !Rd(n)\ } \overline{Ca'} \cup \{Ca_2\}}
$$

This latter rule itself is triggered by the rules

$$
\text{(LLC-Miss)}
$$
$$
\frac{last(caid) = true}{(caid \bullet M \bullet \boxed{dst + \textbf{fetch}(n)}) \xrightarrow{\ !Rd(n)\ } (caid \bullet M \bullet \boxed{dst + \textbf{fetchBl}(n)})}
$$

and

$$
\text{(Flush-One-Line)} \qquad\qquad \text{(Ignore-Flush-One-Line)}
$$
$$
\frac{status(M,n) = mo}{caid \bullet M \bullet \boxed{dst} \xrightarrow{\ ?Rd(n)\ } caid \bullet M \bullet \boxed{dst + \textbf{flush}(n)}} \qquad \frac{status(M,n) \neq mo}{caid \bullet M \bullet dst \xrightarrow{\ ?Rd(n)\ } caid \bullet M \bullet dst}
$$

Together, these rules capture the broadcast mechanism for getting the most recent shared copy of a memory block in the multicore memory system. Rule LLC-Miss corresponds to the case when a last level cache needs to fetch a memory block n, by broadcasting a read message, $!Rd(n)$, to all other caches. This is achieved through triggering the global synchronisation rules Synch and Synch-Dist. While the former identifies the core CR that broadcasts the read message, the latter propagates the message $!Rd(n)$ to the other caches. Depending on the status of block n, the recipient cache will either flush the local modified copy to the main memory (Flush-One-Line) or ignore the message (Ignore-Flush-One-Line).

Clearly, the receiveRd method is simulated by the above rules Flush-One-Line and Ignore-Flush-One-Line. Since we may assume (as argued in Section 2) that the execution of the broadcast method only consists of an interleaving of the processes that are generated by it, it is easy to statically verify that the execution of **this**.broadcast(**this**,n);**this**!fetchBl(n) is simulated by the above rules.

What remains is the correctness of the methods fecthBl, fetchW, and flush (see Figures 16, 17, and 18). This can be established in a similar manner as the above correctness proof for the fetch method. Below we discuss these correctness proofs, omitting details which are straightforward to check.

The condition of the annotation at the beginning of the fetchBl method identifies the path which terminates after the call of the fetchW method (Line 23). It is straightforward to check that this path is simulated by the

rule

$$(\textsc{FetchBl}_3)$$

$$\frac{last(caid) = true \quad select(M,n) = n' \quad n' \neq n \quad status(M,n') = mo}{(caid \bullet M \bullet \boxed{dst + \textbf{fetchBl}(n)}) \rightarrow (caid \bullet M \bullet \boxed{dst + \textbf{flush}(n') + \textbf{fetchW}(n,n')})}$$

which handles the case where a last level cache is blocked on n and the location where n is to be placed is occupied by a modified memory block. To free the location for n, the modified block needs to be flushed to the main memory by the **flush** instruction, and the fetching will continue with the **fetchW** instruction. By definition, no rules correspond to execution paths starting from the beginning of the fetchBl method and leading to the synchronous calls of the getStatus method (Lines 7, 17, and 25).

The condition of the annotation status != Nothing : LLC-Fetch-Unblock identifies the path which leads from the execution of the getStatus method of the next level cache (Line 7) and subsequent execution of the **else**-branch of the subsequent **if**-statement to termination. It is straightforward to check that this path is simulated by the rule

$$(\textsc{LC-Fetch-Unblock})$$

$$\frac{lid(caid') = lid(caid) + 1 \quad cid(caid) = cid(caid') \quad n \in dom(M')}{\begin{array}{c}(caid \bullet M \bullet \boxed{dst + \textbf{fetchBl}(n)}), (caid' \bullet M' \bullet dst') \rightarrow \\ (caid \bullet M \bullet \boxed{dst + \textbf{fetch}(n)}), (caid' \bullet M' \bullet dst')\end{array}}$$

which corresponds to the case where n, i.e., the memory block that the cache is blocked on, is found in the next level cache. The rule unblocks the cache by trying to fetch n again (replacing **fetchBl**(n) with **fetch**(n) in dst). By definition, no rule corresponds to the path which leads to termination after execution of the **then**-branch. In this case the execution of the fetchBl method involves busy waiting for the status of the given address in the next level cache to become defined.

The annotation $\textsc{FetchBl}_1$ identifies the path which leads from execution of the getStatus method of the main memory and the update cacheMemory = put(cacheMemory,n,status) (Line 17 to termination. It is straightforward to check that this path is simulated by the rule

$$(\textsc{FetchBl}_1)$$

$$\frac{last(caid) = true \quad select(M,n) = n \quad s = status(\overline{M}, n)}{(caid \bullet \boxed{M \bullet dst + \textbf{fetchBl}(n)}), \overline{M} \rightarrow (caid \bullet \boxed{M'[n \mapsto s] \bullet dst}), \overline{M}}$$

which unblocks a last level cache which has space to fetch the memory block n from main memory $\overline{M}$. Similarly, the annotation $\textsc{FetchBl}_2$ identifies the path which leads from execution of the getStatus method of the main memory (Line 25) and the subsequent updates removeKey(cacheMemory,selected_n) and cacheMemory = put(cacheMemory,n,status) to termination. It is straightforward to check that this path is simulated by the rule

$$(\textsc{FetchBl}_2)$$

$$\frac{last(caid) = true \quad select(M,n) = n' \quad n' \neq n \quad status(M,n') \neq mo \quad s = status(\overline{M}, n)}{(caid \bullet \boxed{M' \bullet dst + \textbf{fetchBl}(n)}), \overline{M} \rightarrow (caid \bullet \boxed{M'[n' \mapsto \bot, n \mapsto s] \bullet dst}), \overline{M}}$$

which corresponds to a last level cache fetching the memory block n from main memory, but the location where n is to be placed is occupied by a non-modified block. The rule then removes the occupying block and places n into the cache.

Concerning the fetchW method (Figure 17), it is straightforward to check that the path which leads from the **await** statement to its termination is simulated by the rule.

$$(\textsc{FetchW})$$

$$\frac{last(caid) = true \quad status(M,n') \neq mo}{(caid \bullet M \bullet \boxed{dst + \textbf{fetchW}(n,n')}) \rightarrow (caid \bullet M \bullet \boxed{dst + \textbf{fetchBl}(n)})}$$

which handles the case where a last level cache, which fails to replace n' with n because n' was a modified block, can try to fetch n again.

```
1   Unit fetchBl(Address n){
2     // nextLevel==Nothing &
3     // fst(select(cacheMemory, n)) !=n &
4     // select(cacheMemory, n)==Pair(_,Mo): FETCHBL₃;
5     switch (nextLevel) {
6       Just(nextCache) => {
7         Maybe<Status> status = nextCache.getStatus(n);
8         // status!=Nothing: LC-FETCH-UNBLOCK
9         if (status == Nothing){
10          this!fetchBl(n); }
11        else {
12          this!fetch(n); }
13        }
14      _ => {
15        Pair<Address,Status> selected = select(cacheMemory, n);
16        if (fst(selected)==n){
17         Status status = mainMemory.getStatus(n); // FETCHBL₁
18          cacheMemory = put(cacheMemory,n,status); }
19        else {
20          switch (selected) {
21            Pair(selected_n,Mo) => {
22              this!flush(selected_n);
23              this!fetchW(n,selected_n); }
24            Pair(selected_n,_) => {
25              Status status = mainMemory.getStatus(n); // FETCHBL₂
26              cacheMemory = removeKey(cacheMemory,selected_n);
27              cacheMemory = put(cacheMemory,n,status); }
28          } } } } }
```

Fig. 16. The annotated fetchBl method.

```
1     Unit fetchW(Address n,Address n_){
2       await (lookupDefault(cacheMemory,n_, In)!=Mo); // FETCHW
3       this!fetchBl(n);
4       }
```

Fig. 17. The annotated fetchW method.

```
1   Unit flush(Address n) {
2   // lookup(cacheMemory,n)!=Mo: FLUSH₂;
3     switch (lookup(cacheMemory,n)) {
4       Just(Mo) => {
5         mainMemory.setStatus(n,Sh); // FLUSH₁
6         cacheMemory = put(cacheMemory,n,Sh); }
7       _ => skip;
8     } }
```

Fig. 18. The annotated flush method.

The path identified by the condition of the annotation lookup(cacheMemory,n) != Mo : FLUSH₂ of the flush method (Figure 18) is clearly simulated by the rule

$$(\textsc{Flush}_2)$$
$$\frac{status(M,n) \neq mo}{(caid \bullet M \bullet \boxed{dst + \mathbf{flush}(n)}) \rightarrow (caid \bullet M \bullet \boxed{dst})}$$

which just ignores a **flush** instruction if the block is not modified. By definition of this annotation, no rule corresponds with the path leading to the synchronous call of the method setStatus of mainMemory. According to the annotation $\textsc{Flush}_1$, the path which leads from execution of the setStatus method of the main memory to termination, corresponds to the rule

$$
\frac{(\textsc{Flush}_1)}{status(M, n) = mo}
$$

$$
(caid \bullet \; M \bullet dst + \textbf{flush}(n) \;), \; \overline{M} \; \to (caid \bullet \; M[n \mapsto sh] \bullet dst \;), \; \overline{M}[n \mapsto sh]
$$

which flushes a modified memory block n by updating its status to shared in both the cache and the main memory. Clearly, by the simulation relation, the condition of this rule is satisfied in $\alpha(G)$. Further, the update cacheMemory = put(cacheMemory,n,Sh) is simulated by $M[n \mapsto sh]$ and the update mainMemory.setStatus(n,Sh) is simulated by $\overline{M}[n \mapsto sh]$.

6.1   Bisimulation

We briefly discuss how to extend Theorem 2 to a bisimulation between the transitive, reflexive closure of the transition relation $\Rightarrow$ of the ABS multicore program and that of the transition relation $\to$ of the multicore TSS. Such a bisimulation relation then allows to prove both *safety* and *liveness* properties of the ABS multicore program in terms of the multicore TSS. The following theorem states that the multicore TSS is simulated by the ABS program.

Theorem 3. *Let G be a reachable stable global configuration of the ABS multicore program. If $\alpha(G) \to cf$ then there exists a stable global configuration G' such that $\alpha(G') = cf$ and $G \Rightarrow^* G'$.*

Proof. We sketch a proof of this theorem which is based on the correctness of the annotations as established in the proof of Theorem 2. The global structure of the proof of Theorem 3 however involves an analysis of the individual TSS rules. All these rules are triggered by a *dst* instruction. For each such instruction we check statically for each stable point of the corresponding ABS method whether there exists a path to another stable point (*not necessarily the next one*) execution of which corresponds to the TSS rule application.

As an example of this scheme we give an analysis of an application of the rule

$$
\frac{(\textsc{LC-Miss})}{lid(caid') = lid(caid) + 1 \quad cid(caid) = cid(caid') \quad status(M', n) \in \{inv, \bot\}}
$$

$$
(caid \bullet M \bullet \; dst + \textbf{fetch}(n) \;), \; (caid' \bullet \; M' \bullet dst' \;) \to
$$

$$
(caid \bullet M \bullet \; dst + \textbf{fetchBl}(n) \;), \; (caid' \bullet \; M'[n \mapsto \bot] \bullet dst' + \textbf{fetch}(n) \;)
$$

in $\alpha(G)$, triggered by the *dst* instruction **fetch**(n). By definition of α there exists a process instance (either executing or suspended) of the fetch method with the address n as the value of the its formal parameter. Further, the status of the address n of the next level cache, denoted by the nextCache field of the cache object to which this process belongs, is undefined or invalid. We have the following straightforward analysis of the stable points of this process.

In case of the initial stable point and the stable point associated with the call of the remove_inv method (Line 5 of the fetch method, Figure 13), by definition of α there exists a computation $G \Rightarrow^* G'$ which involves in both cases the execution of the path from Line 5 to Line 8 As argued in the proof of Theorem 2, this path corresponds to an application of the LC-Miss rule. Note that execution of the path from Line 1 to Line 5 corresponds to a silent transition in the multicore TSS,

In case of the stable point associated with the call of the swap method (Line 10), by definition of α there exists a computation $G \Rightarrow^* G'$ which involves execution of the path which leads from the return of the swap method via the **else**-branch of the subsequent **if**-statement (note that s == Nothing) to termination of the fetch method, *followed* by execution of the path from the initial stable point of the *newly* generated process instance of the fetch method to Line 8. As argued above, this latter path corresponds to an application of the LC-Miss rule. Execution of the first path corresponds to a silent transition in the multicore TSS (as argued in the proof of Theorem 2).

7 Related Work

There is in general a significant gap between a transition specification and its implementation in a (high-level) parallel programming language [49]. Transition system specifications [46] succinctly formalize operational models and are well-suited for proofs, but direct implementations of such specifications quickly lead to very inefficient implementations. Executable semantic frameworks such as Redex [24], rewriting logic [43,42], and $\mathbb{K}$ [47] reduce this gap, and have been used to develop executable formal models of complex languages like C [23] and Java [12]. The relationship between transition system specifications and rewriting logic semantics has been studied [50] without proposing a general solution for synchronization by label matching. Bijo et al. implemented their multicore memory model [8] in the rewriting logic system Maude [15] using an orchestrator for label matching, but do not provide a correctness proof wrt. the transition system specification. Different semantic styles can be modeled and related inside one framework; for example, the correctness of distributed implementations of KLAIM systems in terms of simulation relations have been studied in rewriting logic [22]. Compared to these works on semantics, we developed a general methodology for proving the correctness of parallel implementations of transition system specifications in the active object language ABS. Our methodology features a new integration of these two formalisms which consists of a formal scheme for annotating ABS programs with transition rules. These annotations provide a high-level specification of the proof obligations for establishing the simulation relation between a transition system specification and its ABS implementation.

Correctness-preserving compilation and refinement is related to correctness proofs for implementations, and ensures that the low-level representation of a program preserves the properties of the high-level model. Examples of this line of work include the B-method [1], which is based on refinement between abstract state machines, type-preserving translations into typed assembly languages [45], and formally verified compilers [36,37], which proves the semantic preservation of a compiler from C to assembler code, but leaves shared-variable concurrency for future work. In contrast to these works our work specifically targets the correctness of parallel systems.

Simulation tools for cache coherence protocols can evaluate performance and efficiency on different architectures (e.g., gems [41] and gem5 [11]). These tools perform evaluations of, e.g., the cache hit/miss ratio and response time, by running benchmark programs written as low-level read and write instructions to memory. Advanced simulators such as Graphite [44] and Sniper [14] run programs on distributed clusters to simulate executions on multicore architectures with thousands of cores. Unlike our work, these simulators are not based on a formal semantics and correctness proofs. Our work complements these simulators by supporting the executable exploration of design choices from a programmer perspective rather from hardware design. Compared to worst-case response time analysis for concurrent programs on multicore architectures [38], our focus is on the underlying data movement rather than the response time.

8 Conclusion

We have introduced in this paper a methodology for proving the correctness of parallel implementations of high-level transition system specifications in the active object language ABS. The proof method consists of establishing a simulation relation between the transition system describing the semantics of the ABS program and the transition system described by the specification. The proof method exploits a general global confluence property of the ABS semantics which allows to abstract from the interleaving of parallel processes and focus on the static analysis of sequential code in the simulation proof. A promising further formalization and tool-supported automation of our methodology is the *symbolic execution* of sequential ABS code in establishing the simulation relation between the ABS program and its specification.

A concern that often arises in parallel execution is fairness: the degree of variability when distributing the computing resources among different parallel components — here, the simulated cores. Fairness of parallel execution can affect the simulation's accuracy in approximating the intended (or idealized) manycore hardware. To ensure fairness of the simulation, we make use of *deployment components* [31] in ABS.

A *Deployment Component* (DC) is an ABS execution location that is created with a number of virtual resources (e.g., execution speed, memory use, network bandwidth), which are shared among its deployed objects. Any annotated statement [Cost: x] S decrements by x the resources of its DC and then completes, or it will stall its computation if there are currently not enough resources remaining; the statement S may continue on the next passage of the global symbolic time where all the resources of the DCs have been renewed, and will eventually complete when its Cost has reached zero.

We make use of this resource modeling of ABS to assign equal (fair) resources of virtual execution speed to the simulated cores of the system. Each Core object is deployed onto a separate DC with fixed Speed(1)

resources. The processing of each instruction has the same cost [Cost: 1] — a generalization, since common processor architectures execute different instructions in different speeds (cycles per instruction); e.g., JUMP is faster than LOAD. The result is that all Cores can execute maximum one instruction in every time interval of the global symbolic clock, and thus no Core can get too far ahead with processing its own instructions — a problem that manifests upon the parallel simulation of N number of cores using a physical machine of M cores, where N is vastly greater than M.

We plan further development of this extension of the ABS multicore model with deployment components for simulating the execution of (object-oriented) programs on multicore architectures. A first such development concerns an extension of the abstract memory model with data. In particular, having the addresses of the memory locations themselves as data allows to model and simulate different data layouts of the dynamically generated object structures.

References

1. Jean-Raymond Abrial. *Modeling in Event-B - System and Software Engineering*. Cambridge University Press, 2010.
2. Elvira Albert, Puri Arenas, Antonio Flores-Montoya, Samir Genaim, Miguel Gómez-Zamalloa, Enrique Martin-Martin, German Puebla, and Guillermo Román-Díez. SACO: static analyzer for concurrent objects. In Erika Ábrahám and Klaus Havelund, editors, *Proceedings of the 20th International Conference on Tools and Algorithms for the Construction and Analysis of Systems (TACAS 2014)*, volume 8413 of *Lecture Notes in Computer Science*, pages 562–567. Springer, 2014.
3. Elvira Albert, Frank S. de Boer, Reiner Hähnle, Einar Broch Johnsen, Rudolf Schlatte, Silvia Lizeth Tapia Tarifa, and Peter Y. H. Wong. Formal modeling and analysis of resource management for cloud architectures: an industrial case study using Real-Time ABS. *Service Oriented Computing and Applications*, 8(4):323–339, 2014.
4. Elvira Albert, Miguel Gómez-Zamalloa, and Miguel Isabel. SYCO: a systematic testing tool for concurrent objects. In Ayal Zaks and Manuel V. Hermenegildo, editors, *Proceedings of the 25th International Conference on Compiler Construction (CC 2016)*, pages 269–270. ACM, 2016.
5. Gregory R. Andrews. *Foundations of Multithreaded, Parallel, and Distributed Programming*. Addison Wesley, 2000.
6. Nikolaos Bezirgiannis, Frank S. de Boer, and Stijn de Gouw. Human-in-the-loop simulation of cloud services. In Flavio De Paoli, Stefan Schulte, and Einar Broch Johnsen, editors, *Proceedings of the 6th IFIP WG 2.14 European Conference on Service-Oriented and Cloud Computing (ESOCC 2017)*, volume 10465 of *Lecture Notes in Computer Science*, pages 143–158. Springer, 2017.
7. Nikolaos Bezirgiannis, Frank S. de Boer, Einar Broch Johnsen, Ka I Pun, and Silvia Lizeth Tapia Tarifa. Implementing SOS with active objects: A case study of a multicore memory system. In Reiner Hähnle and Wil M. P. van der Aalst, editors, *Proceedings of the 22nd International Conference on Fundamental Approaches to Software Engineering (FASE 2019)*, volume 11424 of *Lecture Notes in Computer Science*, pages 332–350. Springer, 2019.
8. Shiji Bijo, Einar Broch Johnsen, Ka I Pun, and S. Lizeth Tapia Tarifa. A Maude framework for cache coherent multicore architectures. In *Proceedings of the 11th International Workshop on Rewriting Logic and Its Applications (WRLA 2016)*, volume 9942 of *Lecture Notes in Computer Science*, pages 47–63. Springer, 2016.
9. Shiji Bijo, Einar Broch Johnsen, Ka I Pun, and S. Lizeth Tapia Tarifa. A formal model of parallel execution on multicore architectures with multilevel caches. In *Proceedings of the 14th International Conference on Formal Aspects of Component Software (FACS 2017)*, volume 10487 of *Lecture Notes in Computer Science*, pages 58–77. Springer, 2017.
10. Shiji Bijo, Einar Broch Johnsen, Ka I Pun, and Silvia Lizeth Tapia Tarifa. A formal model of data access for multicore architectures with multilevel caches. *Science of Computer Programming*, 179:24–53, 2019.
11. Nathan Binkert, Bradford Beckmann, Gabriel Black, Steven K. Reinhardt, Ali Saidi, Arkaprava Basu, Joel Hestness, Derek R. Hower, Tushar Krishna, Somayeh Sardashti, Rathijit Sen, Korey Sewell, Muhammad Shoaib, Nilay Vaish, Mark D. Hill, and David A. Wood. The gem5 simulator. *SIGARCH Computer Architecture News*, 39(2):1–7, 2011.
12. Denis Bogdanas and Grigore Rosu. K-Java: A complete semantics of java. In Sriram K. Rajamani and David Walker, editors, *Proc. 42nd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL 2015)*, pages 445–456. ACM, 2015.
13. Roland N. Bol and Jan Friso Groote. The meaning of negative premises in transition system specifications. *J. ACM*, 43(5):863–914, 1996.
14. Trevor E. Carlson, Wim Heirman, and Lieven Eeckhout. Sniper: Exploring the level of abstraction for scalable and accurate parallel multi-core simulation. In *Proceedings of International Conference for High Performance Computing, Networking, Storage and Analysis (SC)*, pages 52:1–52:12. ACM, 2011.
15. Manuel Clavel, Francisco Durán, Steven Eker, Patrick Lincoln, Narciso Martí-Oliet, José Meseguer, and Carolyn L. Talcott, editors. *All About Maude - A High-Performance Logical Framework, How to Specify, Program and Verify Systems in Rewriting Logic*, volume 4350 of *Lecture Notes in Computer Science*. Springer, 2007.
16. David E. Culler, Anoop Gupta, and Jaswinder Pal Singh. *Parallel Computer Architecture: A Hardware/Software Approach*. Morgan Kaufmann Publishers Inc., 1st edition, 1997.

17. Ferruccio Damiani, Reiner Hähnle, Eduard Kamburjan, and Michael Lienhardt. A unified and formal programming model for deltas and traits. In Marieke Huisman and Julia Rubin, editors, *Proceedings of the 20th International Conference on Fundamental Approaches to Software Engineering (FASE 2017)*, volume 10202 of *Lecture Notes in Computer Science*, pages 424–441. Springer, 2017.

18. Frank de Boer, Vlad Serbanescu, Reiner Hähnle, Ludovic Henrio, Justine Rochas, Crystal Chang Din, Einar Broch Johnsen, Marjan Sirjani, Ehsan Khamespanah, Kiko Fernandez-Reyes, and Albert Mingkun Yang. A survey of active object languages. *ACM Comput. Surv.*, 50(5):76:1–76:39, October 2017.

19. Stijn de Gouw, Frank S. de Boer, Richard Bubel, Reiner Hähnle, Jurriaan Rot, and Dominic Steinhöfel. Verifying OpenJDK's sort method for generic collections. *J. Autom. Reason.*, 62(1):93–126, 2019.

20. Stijn de Gouw, Jacopo Mauro, Behrooz Nobakht, and Gianluigi Zavattaro. Declarative elasticity in ABS. In Marco Aiello, Einar Broch Johnsen, Schahram Dustdar, and Ilche Georgievski, editors, *Proceedings of the 5th IFIP WG 2.14 European Conference Service-Oriented and Cloud Computing (ESOCC 2016)*, volume 9846 of *Lecture Notes in Computer Science*, pages 118–134. Springer, 2016.

21. Crystal Chang Din, Richard Bubel, and Reiner Hähnle. KeY-ABS: A deductive verification tool for the concurrent modelling language ABS. In *CADE*, volume 9195 of *Lecture Notes in Computer Science*, pages 517–526. Springer, 2015.

22. Jonas Eckhardt, Tobias Mühlbauer, José Meseguer, and Martin Wirsing. Semantics, distributed implementation, and formal analysis of KLAIM models in Maude. *Sci. Comput. Program.*, 99:24–74, 2015.

23. Chucky Ellison and Grigore Rosu. An executable formal semantics of C with applications. In John Field and Michael Hicks, editors, *Proceedings of the 39th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL 2012)*, pages 533–544. ACM, 2012.

24. Matthias Felleisen, Robert Bruce Findler, and Matthew Flatt. *Semantics Engineering with PLT Redex*. The MIT Press, 2009.

25. Rob van Glabbeek. On the meaning of transition system specifications. In Jorge A. Pérez and Jurriaan Rot, editors, *Proceedings Combined 26th International Workshop on Expressiveness in Concurrency and 16th Workshop on Structural Operational Semantics (EXPRESS/SOS 2019)*, volume 300 of *EPTCS*, pages 69–85, 2019.

26. Jan Friso Groote and Frits W. Vaandrager. Structured operational semantics and bisimulation as a congruence. *Inf. Comput.*, 100(2):202–260, 1992.

27. Reiner Hähnle. The abstract behavioral specification language: A tutorial introduction. In *FMCO 2012*, volume 7866 of *Lecture Notes in Computer Science*, pages 1–37. Springer, 2012.

28. Carl Hewitt, Peter Bishop, and Richard Steiger. A Universal Modular ACTOR Formalism for Artificial Intelligence. In *Proceedings of the 3rd International Joint Conference on Artificial Intelligence*, IJCAI'73, pages 235–245, San Francisco, CA, USA, 1973. Morgan Kaufmann Publishers Inc.

29. Einar Broch Johnsen, Reiner Hähnle, Jan Schäfer, Rudolf Schlatte, and Martin Steffen. ABS: A core language for abstract behavioral specification. In Bernhard Aichernig, Frank S. de Boer, and Marcello M. Bonsangue, editors, *Proc. 9th International Symposium on Formal Methods for Components and Objects (FMCO 2010)*, volume 6957 of *Lecture Notes in Computer Science*, pages 142–164. Springer, 2011.

30. Einar Broch Johnsen, Jia-Chun Lin, and Ingrid Chieh Yu. Comparing AWS deployments using model-based predictions. In *ISoLA (2)*, volume 9953 of *Lecture Notes in Computer Science*, pages 482–496, 2016.

31. Einar Broch Johnsen, Rudolf Schlatte, and S.Ĺizeth Tapia Tarifa. Integrating deployment architectures and resource consumption in timed object-oriented models. *Journal of Logical and Algebraic Methods in Programming*, 84(1):67–91, 2015.

32. Eduard Kamburjan. Detecting deadlocks in formal system models with condition synchronization. *Electron. Commun. Eur. Assoc. Softw. Sci. Technol.*, 76, 2018.

33. Eduard Kamburjan, Reiner Hähnle, and Sebastian Schön. Formal modeling and analysis of railway operations with active objects. *Sci. Comput. Program.*, 166:167–193, 2018.

34. Eduard Kamburjan, Marco Scaletta, and Nils Rollshausen. Crowbar: Behavioral symbolic execution for deductive verification of active objects. *CoRR*, abs/2102.10127, 2021.

35. Eduard Kamburjan and Jonas Stromberg. Tool support for validation of formal system models: Interactive visualization and requirements traceability. In *F-IDE@FM*, volume 310 of *EPTCS*, pages 70–85, 2019.

36. Xavier Leroy. Formal verification of a realistic compiler. *Commun. ACM*, 52(7):107–115, 2009.

37. Xavier Leroy. A formally verified compiler back-end. *J. Autom. Reasoning*, 43(4):363–446, 2009.

38. Yan Li, Vivy Suhendra, Yun Liang, Tulika Mitra, and Abhik Roychoudhury. Timing analysis of concurrent programs running on shared cache multi-cores. In Theodore P. Baker, editor, *Proceedings of the 30th IEEE Real-Time Systems Symposium (RTSS 2009)*, pages 57–67. IEEE Computer Society, 2009.

39. Jia-Chun Lin, Jacopo Mauro, Thomas Brox Røst, and Ingrid Chieh Yu. A model-based scalability optimization methodology for cloud applications. In *Proceedings of the 7th International Symposium on Cloud and Service Computing (SC^2 2017)*, pages 163–170. IEEE Computer Society, 2017.

40. Jia-Chun Lin, Ingrid Chieh Yu, Einar Broch Johnsen, and Ming-Chang Lee. ABS-YARN: A formal framework for modeling hadoop YARN clusters. In Perdita Stevens and Andrzej Wasowski, editors, *Proceedings of the 19th International Conference on Fundamental Approaches to Software Engineering (FASE 2016)*, volume 9633 of *Lecture Notes in Computer Science*, pages 49–65. Springer, 2016.

41. Milo M. K. Martin, Daniel J. Sorin, Bradford M. Beckmann, Michael R. Marty, Min Xu, Alaa R. Alameldeen, Kevin E. Moore, Mark D. Hill, and David A. Wood. Multifacet's general execution-driven multiprocessor simulator (GEMS) toolset. *SIGARCH Computer Architecture News*, 33(4):92–99, 2005.
42. José Meseguer and Grigore Rosu. The rewriting logic semantics project. *Theor. Comput. Sci.*, 373(3):213–237, 2007.
43. José Meseguer and Grigore Rosu. The rewriting logic semantics project: A progress report. *Inf. Comput.*, 231:38–69, 2013.
44. Jason E. Miller, Harshad Kasture, George Kurian, Charles Gruenwald, Nathan Beckmann, Christopher Celio, Jonathan Eastep, and Anant Agarwal. Graphite: A distributed parallel simulator for multicores. In *Proceedings of the 16th International Symposium on High-Performance Computer Architecture (HPCA)*, pages 1–12. IEEE Computer Society, 2010.
45. J. Gregory Morrisett, David Walker, Karl Crary, and Neal Glew. From system F to typed assembly language. *ACM Trans. Program. Lang. Syst.*, 21(3):527–568, 1999.
46. Gordon D. Plotkin. A structural approach to operational semantics. *Journal of Logic and Algebraic Programming*, 60-61:17–139, 2004.
47. Grigore Rosu. 𝕂: A semantic framework for programming languages and formal analysis tools. In *Dependable Software Systems Engineering*, pages 186–206. IOS Press, 2017.
48. Rudolf Schlatte, Einar Broch Johnsen, Eduard Kamburjan, and Silvia Lizeth Tapia Tarifa. Modeling and analyzing resource-sensitive actors: A tutorial introduction. In *COORDINATION 2021*, volume 12717 of *Lecture Notes in Computer Science*, pages 3–19. Springer, 2021.
49. Rudolf Schlatte, Einar Broch Johnsen, Jacopo Mauro, Silvia Lizeth Tapia Tarifa, and Ingrid Chieh Yu. Release the beasts: When formal methods meet real world data. In *It's All About Coordination - Essays to Celebrate the Lifelong Scientific Achievements of Farhad Arbab*, volume 10865 of *Lecture Notes in Computer Science*, pages 107–121. Springer, 2018.
50. Traian-Florin Serbanuta, Grigore Rosu, and José Meseguer. A rewriting logic approach to operational semantics. *Inf. Comput.*, 207(2):305–340, 2009.
51. Yan Solihin. *Fundamentals of Parallel Multicore Architecture*. Chapman & Hall/CRC, 1st edition, 2015.
52. Daniel J. Sorin, Mark D. Hill, and David A. Wood. *A Primer on Memory Consistency and Cache Coherence*. Morgan & Claypool Publishers, 1st edition, 2011.
53. Gianluca Turin, Andrea Borgarelli, Simone Donetti, Einar Broch Johnsen, Silvia Lizeth Tapia Tarifa, and Ferruccio Damiani. A formal model of the kubernetes container framework. In Tiziana Margaria and Bernhard Steffen, editors, *Proceedings of the 9th International Symposium on Leveraging Applications of Formal Methods (ISoLA 2020)*, volume 12476 of *Lecture Notes in Computer Science*, pages 558–577. Springer, 2020.
54. Lars Tveito, Einar Broch Johnsen, and Rudolf Schlatte. Global reproducibility through local control for distributed active objects. In Heike Wehrheim and Jordi Cabot, editors, *Proc. 23rd International Conference on Fundamental Approaches to Software Engineering (FASE 2020)*, volume 12076 of *Lecture Notes in Computer Science*, pages 140–160. Springer, 2020.

A   Multicore TSS

The multicore TSS is structured in terms of separate TSS's for the cores, caches, an global synchronization. In general, we assume that the unlabelled transitions which describe the behavior of the individual cores and caches are applied in the context of a configuration *cf*, and we omit the straightforward context rule here. On the other hand, for the labelled transitions we introduce explicit synchronization rules for lifting them to a particular context.

Transition Rules for Cores

Figure 19 shows the transition rules for the basic core instructions **read**(r), **readBl**(r), **write**(r), and **writeBl**(r).

Transition Rules for Caches

These rules are further structured in terms of separate TSS's for the individual *rst* instructions (Figures 20, 21, 22, and 23).

Transition Rules for Global Synchronization

These rules are further structured in terms of a TSS for labelled transitions (Figure 24) and a TSS of rules for matching these labelled transitions (Figure 25).

$$(\text{PRRD}_1)$$
$$\frac{first(caid) = true \quad cid(caid) = c \quad status(M,n) \in \{sh, mo\}}{(c \bullet \text{read}(n); rst \,), \, (caid \bullet M \bullet dst) \to (c \bullet rst \,), \, (caid \bullet M \bullet dst)}$$

$$(\text{PRRD}_2)$$
$$\frac{first(caid) = true \quad cid(caid) = c \quad status(M,n) \in \{inv, \bot\}}{\begin{array}{c}(c \bullet \text{read}(n); rst \,), \, (caid \bullet M \bullet dst \,) \to \\ (c \bullet \text{readBl}(n); rst \,), \, (caid \bullet M[n \mapsto \bot] \bullet dst + \text{fetch}(n) \,)\end{array}}$$

$$(\text{PRRD}_3)$$
$$\frac{first(caid) = true \quad cid(caid) = c \quad n \in dom(M)}{(c \bullet \text{readBl}(n); rst \,), \, (caid \bullet M \bullet dst) \to (c \bullet \text{read}(n); rst \,), \, (caid \bullet M \bullet dst)}$$

$$(\text{PRWR}_1)$$
$$\frac{first(caid) = true \quad cid(caid) = c \quad status(M,n) = mo}{(c \bullet \text{write}(n); rst \,), \, (caid \bullet M \bullet dst) \to (c \bullet rst \,), \, (caid \bullet M \bullet dst)}$$

$$(\text{PRWR}_2)$$
$$\frac{first(caid) = true \quad cid(caid) = c \quad status(M,n) = sh}{(c \bullet \text{write}(n); rst \,), \, (caid \bullet M \bullet dst) \xrightarrow{!RdX(n)} (c \bullet rst \,), \, (caid \bullet M[n \mapsto mo] \bullet dst)}$$

$$(\text{PRWR}_3)$$
$$\frac{first(caid) = true \quad cid(caid) = c \quad status(M,n) \in \{inv, \bot\}}{\begin{array}{c}(c \bullet \text{write}(n); rst \,), \, (caid \bullet M \bullet dst \,) \to \\ (c \bullet \text{writeBl}(n); rst \,), \, (caid \bullet M[n \mapsto \bot] \bullet dst + \text{fetch}(n) \,)\end{array}}$$

$$(\text{PRWR}_4)$$
$$\frac{first(caid) = true \quad cid(caid) = c \quad n \in dom(M)}{(c \bullet \text{writeBl}(n); rst \,), \, (caid \bullet M \bullet dst) \to (c \bullet \text{write}(n); rst \,), \, (caid \bullet M \bullet dst)}$$

Fig. 19. Transition rules for $\text{read}(r)$, $\text{readBl}(r)$, $\text{write}(r)$, and $\text{writeBl}(r)$.

B   Multicore ABS

In this section, we collect all the methods of the ABS models that we have discussed in the paper to provide a full view of the ABS implementation of the multicore memory system.

$$(\text{LC-Hit}_1)$$

$$lid(caid') = lid(caid) + 1 \quad cid(caid) = cid(caid')$$
$$select(M, n) = m \quad n \neq m \quad M(m) = s \quad M'(n) = s' \quad s' = sh \vee s' = mo$$

$$(caid \bullet \boxed{M} \bullet dst + \textbf{fetch}(n)), \ (caid' \bullet \boxed{M'} \bullet dst') \rightarrow$$
$$(caid \bullet \boxed{M[m \mapsto \bot, n \mapsto s']} \bullet dst), \ (caid' \bullet \boxed{M'[n \mapsto \bot, m \mapsto s]} \bullet dst')$$

$$(\text{LC-Hit}_2)$$

$$lid(caid') = lid(caid) + 1 \quad cid(caid) = cid(caid')$$
$$select(M, n) = n \quad M'(n) = s' \quad s' = sh \vee s' = mo$$

$$(caid \bullet \boxed{M} \bullet dst + \textbf{fetch}(n)), \ (caid' \bullet \boxed{M'} \bullet dst') \rightarrow$$
$$(caid_i \bullet \boxed{M[n \mapsto s']} \bullet dst), \ (caid' \bullet \boxed{M'[n \mapsto \bot]} \bullet dst')$$

$$(\text{LC-Miss})$$

$$lid(caid') = lid(caid) + 1 \quad cid(caid) = cid(caid') \quad status(M', n) \in \{inv, \bot\}$$

$$(caid \bullet M \bullet \boxed{dst + \textbf{fetch}(n)}), \ (caid' \bullet \boxed{M' \bullet dst'}) \rightarrow$$
$$(caid \bullet M \bullet \boxed{dst + \textbf{fetchBl}(n)}), \ (caid' \bullet \boxed{M'[n \mapsto \bot] \bullet dst' + \textbf{fetch}(n)})$$

$$(\text{LLC-Miss})$$

$$last(caid) = true$$

$$(caid \bullet M \bullet \boxed{dst + \textbf{fetch}(n)}) \xrightarrow{!Rd(n)} (caid \bullet M \bullet \boxed{dst + \textbf{fetchBl}(n)})$$

Fig. 20. Transition rules for $\textbf{fetch}(n)$.

$$(\text{FetchBl}_1)$$

$$last(caid) = true \quad select(M, n) = n \quad s = status(\overline{M}, n)$$

$$(caid \bullet \boxed{M \bullet dst + \textbf{fetchBl}(n)}), \ \overline{M} \rightarrow (caid \bullet \boxed{M'[n \mapsto s] \bullet dst}), \ \overline{M}$$

$$(\text{FetchBl}_2)$$

$$last(caid) = true \quad select(M, n) = n' \quad n' \neq n \quad status(M, n') \neq mo \quad s = status(\overline{M}, n)$$

$$(caid \bullet \boxed{M' \bullet dst + \textbf{fetchBl}(n)}), \ \overline{M} \rightarrow (caid \bullet \boxed{M'[n' \mapsto \bot, n \mapsto s] \bullet dst}), \ \overline{M}$$

$$(\text{FetchBl}_3)$$

$$last(caid) = true \quad select(M, n) = n' \quad n' \neq n \quad status(M, n') = mo$$

$$(caid \bullet M \bullet \boxed{dst + \textbf{fetchBl}(n)}) \rightarrow (caid \bullet M \bullet \boxed{dst + \textbf{flush}(n') + \textbf{fetchW}(n, n')})$$

$$(\text{LC-Fetch-Unblock})$$

$$lid(caid') = lid(caid) + 1 \quad cid(caid) = cid(caid') \quad n \in dom(M')$$

$$(caid \bullet M \bullet \boxed{dst + \textbf{fetchBl}(n)}), \ (caid' \bullet M' \bullet dst') \rightarrow$$
$$(caid \bullet M \bullet \boxed{dst + \textbf{fetch}(n)}), \ (caid' \bullet M' \bullet dst')$$

Fig. 21. Transition rules for $\textbf{fetchBl}(n)$.

$$(\text{FetchW})$$
$$\dfrac{last(caid) = true \quad status(M, n') \neq mo}{(caid \bullet M \bullet \boxed{dst + \textsf{fetchW}(n, n')}) \rightarrow (caid \bullet M \bullet \boxed{dst + \textsf{fetchBl}(n)})}$$

Fig. 22. Transition rule for $\textsf{fetchW}(n, n')$.

$$(\text{Flush}_1)$$
$$\dfrac{status(M, n) = mo}{(caid \bullet \boxed{M \bullet dst + \textsf{flush}(n)}), \; \boxed{\overline{M}} \rightarrow (caid \bullet \boxed{M[n \mapsto sh] \bullet dst}), \; \boxed{\overline{M}[n \mapsto sh]}}$$

$$(\text{Flush}_2)$$
$$\dfrac{status(M, n) \neq mo}{(caid \bullet M \bullet \boxed{dst + \textsf{flush}(n)}) \rightarrow (caid \bullet M \bullet \boxed{dst})}$$

Fig. 23. Transition rules for $\textsf{flush}(n)$.

$$(\text{Invalidate-One-Line})$$
$$\dfrac{status(M, n) = sh}{caid \bullet \boxed{M} \bullet dst \xrightarrow{?RdX(n)} caid \bullet \boxed{M[n \mapsto inv]} \bullet dst}$$

$$(\text{Ignore-Invalidate-One-Line})$$
$$\dfrac{status(M, n) \in \{inv, \bot\}}{caid \bullet M \bullet dst \xrightarrow{?RdX(n)} caid \bullet M \bullet dst}$$

$$(\text{Flush-One-Line})$$
$$\dfrac{status(M, n) = mo}{caid \bullet M \bullet \boxed{dst} \xrightarrow{?Rd(n)} caid \bullet M \bullet \boxed{dst + \textsf{flush}(n)}}$$

$$(\text{Ignore-Flush-One-Line})$$
$$\dfrac{status(M, n) \neq mo}{caid \bullet M \bullet dst \xrightarrow{?Rd(n)} caid \bullet M \bullet dst}$$

Fig. 24. Labelled input transitions.

$$(\text{Synch-Dist})$$
$$\dfrac{Ca_1 \notin \overline{Ca} \quad \overline{Ca} \xrightarrow{!Rd(n)} \overline{Ca'} \quad Ca_1 \xrightarrow{?Rd(n)} Ca_2'}{\overline{Ca} \cup \{Ca_1\} \xrightarrow{!Rd(n)} \overline{Ca'} \cup \{Ca_2\}}$$

$$(\text{Synch})$$
$$\dfrac{\overline{Ca} \xrightarrow{!Rd(n)} \overline{Ca'}}{\langle \overline{CR}, \overline{Ca}, M \rangle \rightarrow \langle \overline{CR}, \overline{Ca'}, M \rangle}$$

$$(\text{Synch-DistX})$$
$$\dfrac{Ca_1 \notin \overline{Ca} \quad CR, \overline{Ca} \xrightarrow{!RdX(n)} CR', \overline{Ca'} \quad Ca_1 \xrightarrow{?RdX(n)} Ca_2}{CR, \overline{Ca} \cup \{Ca_1\} \xrightarrow{!RdX(n)} CR', \overline{Ca'} \cup \{Ca_2\}}$$

$$(\text{SynchX})$$
$$\dfrac{CR \notin \overline{CR_1} \quad CR, \overline{Ca} \xrightarrow{!RdX(n)} CR', \overline{Ca'}}{\begin{array}{c} \langle \overline{CR_1} \cup \{CR\}, \overline{Ca}, M \rangle \\ \rightarrow \langle \overline{CR_1} \cup \{CR'\}, \overline{Ca'}, M[n \mapsto inv] \rangle \end{array}}$$

Fig. 25. Transition rules for global synchronization/broadcast.

```
1  Unit run() {
2    if currentTask!=Nil {
3    switch (currentTask) {
4       Cons(rst, rest) =>
5       switch (rst) {
6         Read(n) => {
7             removed = l1.remove_inv(n); // removed==True: PRRD₂; removed==False: PRRD₁
8             if (removed){
9               l1!fetch(n);
10              currentTask = Cons(ReadBl(n),rest); }
11            else {currentTask = rest; } }
12         ReadBl(n) => {
13           status = l1.getStatus(n); // status!=Nothing: PRRD₃
14           if (status != Nothing) currentTask = Cons(Read(n),rest); }
15         Write(n) => {
16           status = l1.getStatus(n); // status==Just(Mo): PRWR₁
17           switch (status) {
18             Just(Mo) => {currentTask = rest; }
19             Just(Sh) => {
20               Bool res = l1.broadcastX(n); // res==True: PRWR₂/SYNCHX
21               if (res) {currentTask = rest;  } }
22             _ => { Bool removed = l1.remove_inv(n); // removed==True: PRWR₃
23                 if (removed){l1!fetch(n); currentTask = Cons(WriteBl(n),rest); } } } }
24         WriteBl(n) => {
25           Maybe<Status> status = l1.getStatus(n); // status!=Nothing: PRWR₄
26           if (status != Nothing)
27               currentTask = Cons(Write(n),rest); }
28       }
29    this ! run(); }
30  }
```

Fig. 26. The annotated run method.

```
1   Maybe<Status> getStatus(Address n) { return lookup(cacheMemory,n); }

1   Bool remove_inv(Address n){
2     Bool answer = False;
3     switch (lookup(cacheMemory,n)) {
4       Nothing => { answer = True; }
5       Just(In) =>{ cacheMemory = removeKey(cacheMemory,n); answer = True; }
6       _ => skip; }
7     return answer; }

1   Bool broadcastX(Address n) {
2     Bool res = False;
3     await bus!lock(); //(lookup(cacheMemory,n) ==Just(Sh)): PRWR₂/SYNCH
4     if (lookup(cacheMemory,n) ==Just(Sh)){
5       bus.sendRdX(this, n);
6       cacheMemory = put(cacheMemory,n,Mo);
7       res = True; }
8       bus.release();
9     return res; }
```

Fig. 27. Methods getStatus, remove_inv, and broadcastX of class Cache.

```
1   Unit receiveRdX(Address n,IBarrier start,IBarrier end) {
2     // lookup(cacheMemory,n))==Just(Sh): INVALIDATE-ONE-LINE;
3     // lookup(cacheMemory,n))!=Just(Sh): IGNORE-INVALIDATE-ONE-LINE
4     await start!synchronize();
5     switch (lookup(cacheMemory,n)) {
6       Just(Sh) => {cacheMemory = put(cacheMemory,n,In); }
7       _ => skip;
8     }
9     end.synchronize(); }
```

Fig. 28. The annotated receiveRdX method.

```
1   Unit fetch(Address n){
2   // nextLevel==Nothing: LLC-MISS/SYNCH
3     switch (nextLevel) {
4       Just(nextCache) => {
5         Bool removed = nextCache.remove_inv(n); // removed==true: LC-MISS
6         if (removed){
7           nextCache!fetch(n);
8           this!fetchBl(n); }
9         else { Pair<Address,Status> selected = select(cacheMemory, n);
10          Maybe<Status> s = nextCache.swap(n,selected);
11          // s!=Nothing\& fst(selected)==n: LC-HIT₂;
12          // s!=Nothing\& fst(selected)!=n: LC-HIT₁
13          if (s != Nothing){ if (fst(selected)!=n){
14            cacheMemory = removeKey(cacheMemory,fst(selected)); }
15            cacheMemory = put(cacheMemory, n,fromJust(s)); }
16          else this!fetch(n); } }
17        _ => { this.broadcast(n);
18              this!fetchBl(n); } } }
19
20  Unit broadcast(Address n){
21    await bus!lock();
22    bus.sendRd(this, n);
23    bus.release(); }
```

Fig. 29. The annotated fetch method.

```
1   Maybe<Status> swap(Address n_out, Pair<Address,Status> n_in) {
2     Maybe<Status> tmp = Nothing;
3     switch (lookup(cacheMemory,n_out)) {
4       Nothing => skip;
5       Just(In) => skip;
6       _ => {
7         tmp = lookup(cacheMemory,n_out);
8         cacheMemory = removeKey(cacheMemory,n_out);
9         if (fst(n_in)!=n_out) {
10          cacheMemory = put(cacheMemory, fst(n_in), snd(m_in)); } } }
11    return tmp; }
```

Fig. 30. The swap method.

```
1   Unit receiveRd(Address n,IBarrier start,IBarrier end) {
2     // lookup(cacheMemory,n))==Just(Mo): FLUSH-ONE-LINE;
3     // lookup(cacheMemory,n))!=Just(Mo): IGNORE-FLUSH-ONE-LINE
4     await start!synchronize();
5     switch (lookup(cacheMemory,n)) {
6       Just(Mo) => this!flush(n);
7       _ => skip;
8     }
9     end!synchronize(); }
```

Fig. 31. The annotated receiveRd method.

```
1   Unit fetchBl(Address n){
2     // nextLevel==Nothing &
3     // fst(select(cacheMemory, n)) !=n &
4     // select(cacheMemory, n)==Pair(_ ,Mo): FETCHBL₃;
5     switch (nextLevel) {
6       Just(nextCache) => {
7         Maybe<Status> status = nextCache.getStatus(n);
8         // status!=Nothing: LC-FETCH-UNBLOCK
9         if (status == Nothing){
10          this!fetchBl(n); }
11        else {
12          this!fetch(n); }
13        }
14        _ => {
15        Pair<Address,Status> selected = select(cacheMemory, n);
16        if (fst(selected)==n){
17         Status status = mainMemory.getStatus(n); // FETCHBL₁
18          cacheMemory = put(cacheMemory,n,status); }
19        else {
20          switch (selected) {
21            Pair(selected_n,Mo) => {
22              this!flush(selected_n);
23              this!fetchW(n,selected_n); }
24            Pair(selected_n,_) => {
25              Status status = mainMemory.getStatus(n); // FETCHBL₂
26              cacheMemory = removeKey(cacheMemory,selected_n);
27              cacheMemory = put(cacheMemory,n,status); }
28          } } } } }
```

Fig. 32. The annotated fetchBl method.

```
1     Unit fetchW(Address n,Address n_){
2      await (lookupDefault(cacheMemory,n_, In)!=Mo); // FETCHW
3      this!fetchBl(n);
4      }
```

Fig. 33. The annotated fetchW method.

```
1   Unit flush(Address n) {
2   // lookup(cacheMemory,n)!=Mo: FLUSH₂;
3     switch (lookup(cacheMemory,n)) {
4       Just(Mo) => {
5         mainMemory.setStatus(n,Sh); // FLUSH₁
6         cacheMemory = put(cacheMemory,n,Sh); }
7       _ => skip;
8     } }
```

Fig. 34. The annotated flush method.